

Hedging Market Exposures

VISIT...

LANZAROTE
Caliente.COM

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Finance series contains books written specifically for finance and investment professionals as well as sophisticated individual investors and their financial advisors. Book topics range from portfolio management to e-commerce, risk management, financial engineering, valuation and financial instrument analysis, as well as much more.

For a list of available titles, visit our Web site at www.WileyFinance.com.

Hedging Market Exposures

*Identifying and Managing
Market Risks*

OLEG V. BYCHUK
BRIAN J. HAUGHEY



WILEY

John Wiley & Sons, Inc.

Copyright © 2011 by Oleg Bychuk and Brian Haughey. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Bychuk, Oleg V.

Hedging market exposures : identifying and managing market risks / Oleg V. Bychuk and Brian Haughey.

p. cm. – (Wiley finance series)

Includes index.

ISBN 978-0-470-53506-6; ISBN 978-1-1180-8535-6 (ebk); ISBN 978-1-1180-8536-3 (ebk); ISBN 978-1-1180-8537-0 (ebk)

1. Portfolio management. 2. Risk management. 3. Hedging (Finance)

I. Haughey, Brian, J. II. Title.

HG4529.5.B93 2011

332.64'524—dc22

2011007527

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

Preface	ix
Introduction	xi
About the Authors	xvii
CHAPTER 1	
The Economic Environment	1
1.1 Introduction	1
1.2 Inflation and Unemployment	5
1.3 Central Banks and the Money Supply	6
1.4 The Business Cycle	9
1.5 Predicting the Future?	11
1.6 Economic Indicators	11
CHAPTER 2	
Risk: An Introduction	17
2.1 What Is Risk?	17
2.2 Risks of Financial Instruments	19
2.3 Operational Risk	43
2.4 What Risks Are in Your Portfolio? Hidden Hazards	44
2.5 Hedging Market Risks	47
CHAPTER 3	
Asset Modeling	51
3.1 Asset Value	51
3.2 Financial Models	55
3.3 Valuation Principles	70
3.4 Discount Rates Selection	74
3.5 Cash Flow Projection and Asset Valuation	80
3.6 Stochastic Asset Valuation	82
3.7 The Monte Carlo Method	88
3.8 Stochastic Extrapolation	100

CHAPTER 4**Market Exposures and Factor Sensitivities 103**

4.1	From Valuation to Responses and Sensitivities	103
4.2	Response Matrix and Scenario Grid	105
4.3	Stress-Testing	109
4.4	Sensitivities	110
4.5	Interest Rate Sensitivities: Duration, PV01, Convexity, Key Rate Measures	126
4.6	Numerical Evaluation of Sensitivities	137
4.7	Performance Attribution and Completeness Test	139

CHAPTER 5**Quantifying Portfolio Risks 145**

5.1	The Nature of Risk	145
5.2	Standard Risk Measures	149
5.3	Optimal Hedge Sizing	160
5.4	Tail-Risk Measures	162

CHAPTER 6**The Decision to Hedge 171**

6.1	To Hedge or Not to Hedge?	171
6.2	The Hedging Process	178

CHAPTER 7**Constructing a Hedge 193**

7.1	An Ideal Hedge	193
7.2	A Sample Hedge	194
7.3	Static and Dynamic Hedging	200
7.4	Proxy Hedging	205
7.5	Protection versus Upside	208
7.6	Basis Risk	211
7.7	Unintended Consequences	213
7.8	Hedging Credit Risk	214
7.9	Hedging Prepayment, Redemption, and Other Human Behavior Risks	221
7.10	Execution	223

APPENDIX A**Basics of Probability Theory 227**

APPENDIX B	
Elements of Statistics and Time Series Analysis	247
References	255
Glossary	259
Index	281

Preface

The 2007–2010 financial crisis has highlighted the need to identify and control *all* risks in a financial portfolio, both direct and indirect, and to design and deploy a practical proactive dynamic hedging strategy. Simultaneous adverse moves in stock markets, interest rates, and credit spreads, combined with accelerated company defaults and credit rating downgrades, have devastated many portfolios and, in some cases, wiped out years of gains and caused the demise of storied financial institutions. While some well-publicized losses have resulted from spectacular lapses of common sense and the abandonment of the very basic due diligence principles, most could have been contained by sound risk management.

In this book the authors, who have successfully managed through the turmoil a complex, highly leveraged portfolio of asset-backed securities with significant equity, fixed-income, credit, and other exposures, offer a practical guide to the triad of *identifying*, *quantifying*, and *managing* market risks. These three fundamental elements of the risk-control process have not previously received the attention they deserve. Indeed, the tasks of identifying and quantifying risks and then selecting an optimal hedge, when there is not a “perfect” one, are just as, if not more, important as understanding the intricacies of exotic derivative valuation, but are usually overlooked. This book expounds the three elements of the risk triad by focusing on applied risk *management* as opposed to risk *measurement*.

Like the proverbial wolf in sheep’s clothing, many apparently safe investments may contain hidden, and sometimes unexpected, hazards: hedge counterparty exposure, reliance on an insurance guarantee, and performance triggers. In addition, in many actual situations the market exposures of a portfolio of complex assets can be uncertain and, when identified, the appropriate hedges may be difficult to size and implement. Even the detailed composition of the underlying assets is often unavailable or available only with a time delay of months or even quarters, either due to lack of transparency or because of practical limitations. The authors illustrate these challenges and suggest tools for their analysis and mitigation.

In the wake of the “Great Recession,” there has been a flurry of legislative activity aspiring to devise a regulatory structure that would preclude a

reprise. In the end, one cannot legislate away financial crises any more than one can abolish earthquakes; human society will have to live with both. But just as strict building codes can minimize casualties in a tremor, so sound risk management can limit financial losses.

While this book's primary focus is on market risk, all types of risk (financial, market, credit, counterparty, operational, legal, reputational) are interconnected, and a portfolio, and indeed a firm, manager should be concerned with all of them.

Introduction

PORTFOLIO RISKS

We live in an uncertain world, with one consequence being that there is no such thing as a risk-free investment. No matter where we choose to invest our assets, they will be exposed to financial risk in one form or another. Hoarding our life savings under a mattress, for example, will expose us not only to the risk of loss due to fire or theft, but also to the risk of loss due to inflation, because the purchasing power of our savings will likely decrease over time.

An investor, aware that inflation can erode his wealth, who seeks to earn a positive return by investing in an asset faces other, less obvious, hazards. While an investment that he could make in domestic government securities will pay a return that may, although is by no means guaranteed to, exceed the rate of inflation, so preserving or increasing his purchasing power, there is a very real, even if small in developed economies, chance that a given government will abrogate its obligations and pay none, or only a portion, of the promised return on its securities. Investors eyeing the credit markets in Europe at the end of 2010, for example, were particularly aware of this possibility, with credit default spreads reflecting a great deal of market unease and uncertainty. One lesson of the recent financial crisis is that what may to the unwary appear to be a conservative and secure investment may, in reality, be like the proverbial wolf in sheep's clothing, exposing the investor to very real hazards, such as liquidity, correlation, and contagion risks, and exposure to hedge counterparty creditworthiness, or perhaps to reliance on performance triggers and insurance guarantees.

Every single investment that we can make is exposed to a multitude of hazards, with a myriad of previously unacknowledged factors that expose us to the risk of loss. Historically, what we may refer to as “traditional” risk management largely emphasized diversification as a risk mitigant and focused almost exclusively on credit and interest-rate risk. A typical risk manager would ensure that a portfolio's sector allocation requirements were met, while focusing on the duration and credit quality of a portfolio of fixed-income securities, or on the free cash-flow and interest coverage ratio of an

equity portfolio. The market is, however, a much more complex place than many practitioners previously realized. Our aim in this book is to expose the panoply of hazards that confront a portfolio and to describe how they can be managed. We explore these risks, including hidden hazards such as human behavior risk, and discuss how managers must consider not only the credit, interest-rate and earnings risks of a security, but must also pay close attention to the structure of the instrument, the assumptions underlying any valuation model, and the degree to which the return of the security is likely to be influenced by the behavior of other actors in the market and by economic changes in general.

RISK MANAGEMENT TRIAD

Our approach to risk management and hedging is built around a “risk triad,” consisting of the *identification*, *quantification*, and *management* of risk exposures. As practical investment managers, our guiding mantra is to “think outside the box,” and not to be unduly influenced by conventional wisdom. Even if other participants in the market appear to be ignoring or discounting certain risks, it does not mean that those risks are not dangerous to your financial health. In many cases, the market in general may be oblivious to very real dangers, such as, for example, insurance risk, contagion risk, and liquidity risk. While “Mr. Market” generally reflects the consensus view of risk and return, he is not always efficient or accurate, so the prudent investor should always question the basis for market quotes, seeking to understand what influences prices and whether or not they are susceptible to a dislocation.

The investor should be careful to distinguish between “price” and “value,” should analyze an asset based on its estimated intrinsic cashflows, and should examine the various risks to which those cashflows are subject, before deciding, in light of his willingness and ability to take risk, what proportion of those risks should be hedged, noting that not all risks can be effectively and economically offset, while some risks, in view of the investment objective, are most appropriately retained. An investor should ultimately recognize that his view of an asset’s worth must be greater than a rational seller’s estimate of the value of the same asset, assuming no liquidity constraints and that the same investment opportunities are open to both participants. If this were not to be the case, the seller would seek a higher price than the purchaser would be willing to pay.

While each of the three elements in the risk triad poses its own unique challenges, the quantification of risk is often the most difficult. The *value-at-risk* (VaR) methodology is a widely used tool for estimating potential loss amounts, but one that has several weaknesses that can mislead the unwary.

We discuss alternative and supplementary approaches to quantifying risk exposure and we stress the importance of determining an asset's sensitivity to changes in various market factors, such as index levels, interest rates, volatility, and others.

Without presupposing advanced quantitative background on the part of the reader, we explore these risk sensitivities, often referred to collectively as the “Greeks,” in detail, illustrating how they are derived and how they can be used both individually and as part of a response matrix or sensitivity grid. We show by example why care should be taken when using factors such as a price sensitivity—the *delta*—to compare the risk of more than one asset, and we discuss the difference between absolute and relative deltas. We also distinguish between sensitivities that are *linear*—those first-order relationships, where there is a constant relationship between the risk factor and the change in asset value—and those second-order relationships, sometimes referred to as *gammas*, that are *nonlinear*, where the relationship between the risk factor and the change in asset value increases or decreases as a function of the relative size of the risk factor change.

To aid in the quantification of risk, and its subsequent management, we introduce the concept of an *equivalent variation investment* (EVI), the size of an investment in a market instrument such that the instrument's change in value in response to a small move in a market factor will equal the corresponding change in the asset's value. For a given risk sensitivity the corresponding EVI can be used to size an appropriate hedge—noting, of course, that the ability to hedge requires a counterparty willing to take on the risk exposure that the hedger wishes to offload.

RISK MANAGEMENT IN PRACTICE

Some risks are easy to identify, quantify, and manage. For example, it is relatively easy to hedge the price risk of a large-cap equity portfolio using futures contracts. Other risks can be much more difficult, or practically impossible, to hedge. In many real-world applications, the exposures of a portfolio of complex assets are uncertain, and, when identified, the appropriate hedge may be difficult to size and deploy. Credit risk, for example, is a particularly difficult risk to hedge, one in which an asset manager must decide on the geographic, credit quality, duration, and industry sector characteristics of the exposure and then size and choose an optimal, if imperfect, hedge from a limited selection. For a useful review of various risks confronting a firm the reader may consult, for instance, Pickford and Alexander (2001).

In many situations the detailed composition of an instrument or portfolio's underlying assets is unavailable, or available only with a time delay

of months or even quarters, either due to lack of transparency or because of practical limitations (e.g., the asset base may be simply too big to monitor at the individual investment level). In this case the implementation of the three elements of risk control requires skillful judgment and careful consideration.

A manager may decide, after identifying and quantifying the risk exposures facing his portfolio, that, in view of his ability and willingness to take risk and in light of the availability of hedging solutions, he prefers to hedge only a portion of his portfolio, or to remain unhedged. If choosing to hedge some or all of the exposure, he may choose to hedge with linear, rather than nonlinear instruments, balancing cost, opportunity, and risk. We discuss the advantages and disadvantages of hedging with linear and nonlinear instruments and discuss the technique known as *dynamic hedging*, where a hedge is continually re-evaluated and frequently rebalanced in response to changing market conditions and, often, the passage of time.

Many of the risks we discuss in this book are common to all investment portfolios, although some (e.g., manager risk) are relevant for only specific types of financial instrument. The focus of this volume is on the identification, quantification, and management of those risks faced by portfolios of asset-backed securities (ABS). In practice, however, the collateral underlying an ABS transaction can range from credit-card receivables to motion picture ticket receipts, from fast food franchises to airplane engine lease payments, and from mortgage loan repayments to mutual fund fees—in short, practically any conceivable stream of future cash flows. As a result, ABS portfolios can exhibit a full spectrum of risk exposures; thus, the material in the book is relevant for all asset managers, including those responsible for equity, fixed-income, and alternative investment portfolios.

INTENDED AUDIENCE

While it is true that many firms have dedicated risk control functions and teams of quantitative specialists, today's sophisticated markets and complex investment instruments require that *all* market participants understand the risks they face and ways to control them. Addressing this need, the present book will be of special interest to money managers, investment bankers, and structurers, as well as executives at asset management firms, banks, sovereign wealth funds, endowments, family offices, and hedge funds. Additionally, it will be useful to rating agencies, nonfinancial firms (who usually have substantial market exposures that can encompass interest rate, equity, foreign exchange, and energy risks), private investors, and academics. It will also serve as a helpful introduction for students to the real-world challenges faced by market practitioners.

The authors' primary goal has been to address non-risk specialists; for instance, the book does not require extensive quantitative background. The main purpose is to equip non-specialists with the tools necessary to identify and manage the risks in their portfolios. In addition, the non-technical reader will learn the language of quantitative finance and be able to communicate efficiently and confidently with professional risk managers, traders, "quants," sales staff, and clients. This main objective notwithstanding, the offered volume will also prove beneficial to risk professionals, especially those less familiar with structured products.

As a supplement to this book, readers interested in more quantitative or specialized aspects of risk management should consult the many quality publications available. Without the slightest pretense of exhaustiveness, we cite a few useful titles: Dowd (2002), Fabozzi and Fong (1994), Grant (2004), Jorion (2007), and Taleb (1997). An instructive review of various risks, including nonfinancial, confronting a firm is given by Pickford and Alexander (2001). A timely discussion of the ongoing profound changes in global financial markets and their implications for risk management is presented by El-Erian (2008).

BOOK OUTLINE

Hedging Market Exposures is organized in seven chapters that lead the reader through the risk-management process.

The book begins by introducing the economic environment in Chapter 1 and discusses leading indicators that can signal the direction in which the economy is moving and that are used by practitioners to infer what decisions are likely to be made by the Central Bank. These indicators can offer useful insight into the future levels of interest rates.

Chapter 2 provides a comprehensive survey of the risks facing an investment portfolio, including a discussion of some, such as structure risk and correlation risk, generally not covered in other reviews. It provides an overview of the principles of risk control, paying special attention to identifying portfolio risks, finding suitable proxies for them, and quantifying the risk exposure.

In Chapter 3 the authors introduce the concept of asset value and how it relates to price. The components of a valuation model are discussed, and particular emphasis is placed on the sources of uncertainty that can arise in the model. The discounted cash flow technique is described, and various techniques for cash flow projection are illustrated, along with a discussion of interest rate modeling. The chapter concludes with a detailed review of Monte Carlo modeling.

Chapter 4 opens with a description of how to quantify a portfolio's exposure to various risks through the use of a response matrix and scenario grid. The importance of stress-testing, and the use of synchronized stress tests, is emphasized. A thorough discussion of the "Greeks," or first-, second-, and higher-order sensitivities, is included, and their derivation is explained in a language that does not require advance knowledge of differential calculus. Section 4.5 is devoted to an in-depth discussion of sensitivities to interest rates, while section 4.6 explores some of the practical aspects of the numerical evaluation of risk sensitivities. The chapter concludes by considering attribution analysis, which is of particular importance when attempting to explain the profit and loss (P/L) impact of a set of market moves.

Chapter 5 begins with a consideration of the nature of risk, followed by a discussion of absolute risk measures, such as volatility, Sharpe ratio, and drawdown, and relative risk gauges, such as tracking error, beta, and covariance. Section 5.3 combines the concept of equivalent variation investment (EVI) with the insights into relative risk to derive the optimal hedge size. The chapter concludes with a discussion of *tail-risk measures* and a commentary on their inherent limitation of endeavoring to describe rare events whose empirical observations, by their nature, are few or absent. In addition to the commonly used *value at risk* (VaR), the final section of the chapter introduces *expected tail loss* (ETL) and *median tail loss* (MTL).

The rationale for a hedging decision is examined in Chapter 6, followed by a detailed description of the steps involved in the hedging process. This includes a discussion of how to deal with the practical difficulties of finding a suitable hedge instrument and the logistical elements of trade execution.

Chapter 7 explores the construction of a sample hedge and discusses the relative merits of static and dynamic hedging. Special emphasis is placed on the use of effective delta hedging, and the concept of *delta riding*, or *gamma climbing*, is introduced, whereby the hedger can gain from volatility while maintaining the original degree of hedge protection. The practical importance of proxy hedging is described, and the authors then address the unintended consequences that can arise from the implementation of a hedge. The chapter concludes with treatments of the hedging of credit and human behavior risks.

The appendices review probability theory and basic statistics and direct the reader to further resources.

About the Authors

Oleg Bychuk has 11 years of capital markets experience. This includes roles as Head of Risk Management at Julius Baer Investment Management and Head of Risk Management and Quantitative Research at Alternative Asset Managers. He has also held various positions at Citigroup Global Markets, OppenheimerFunds, and Deutsche Bank. Dr. Bychuk holds degrees from Columbia University (PhD) and Lomonosov Moscow State University and has published numerous articles.

Brian Haughey is Director of the Investment Center at Marist College, in Poughkeepsie, NY, where he teaches Finance and oversees the Student Managed Investment Fund. Previously, he worked in Citigroup for more than 10 years, most recently in the Global Special Situations Group, where he headed Citigroup's Mutual Fund Fee business. Prior to that he was employed at Fitch Ratings and at various institutions in Europe. Mr. Haughey holds an MBS degree from University College Dublin and the Chartered Financial Analyst and Financial Risk Manager designations.

Hedging Market Exposures

CHAPTER 1

The Economic Environment

*Time present and time past
Are both perhaps present in time future,
And time future contained in time past.*

T.S. Eliot, *Burnt Norton*

1.1 INTRODUCTION

A prerequisite to the successful management and hedging of financial portfolios is an understanding of how those portfolios are valued. As we shall see in Chapter 3, it is useful to distinguish between value and price. In this book we shall refer to the *value* of a financial instrument as the present value of any expected cashflows associated with it; generally, we shall define value as an *intrinsic* measure of the worth of an asset, the value calculated being predicated on assumptions about the cashflows that will accrue to the owner of the asset. A consequence of this categorization is that we can calculate a value for almost any asset with expected cashflows using only two tools: cashflow discounting and statistics. We shall consider *price* to be the *extrinsic* measure of worth that the market places on an asset, whether that price is set by the owner of the asset—the *ask price*—or a prospective purchaser of the asset—the *bid price*.

For many assets, the difference between the market price and the value calculated by market participants will be small. At times, it can be inconvenient or cumbersome to value an asset directly; instead, a price can be determined by comparing the asset to another with generally similar characteristics (e.g., using *matrix pricing* to price a selection of bonds). In some cases, a purchaser may be prepared to pay for an asset a market price that

exceeds the intrinsic value that he¹ calculates; in this case the purchaser can be considered, from his perspective, to be paying a *premium* for the asset. Conversely, a purchaser may be able to purchase the asset at a price lower than the value he assigns to it; in this case, the asset is purchased at what he would consider to be a *discount*. In times of market dislocation, there can be a significant disconnect between value and price. How market sentiment affects this dichotomy is explored in Chapter 2.

It is important that an investor understand the distinction between value and price in order to manage the risks inherent in his portfolio. For some assets, the only cashflows that an owner can expect to receive are the proceeds of their sale. For an asset of this kind, intrinsic value derives wholly from the existence of a subsequent purchaser and the price he is prepared to pay. If an investor is considering purchasing an asset to be held as an investment, or to be used as a medium of exchange, he should consider carefully the propositions upon which the value ascribed to it is based and whether those propositions are likely to persist. If precious stones intended to be used as a store of value have no other utility than as a form of scrip, an investor holding them may be sorely disappointed when seeking to subsequently dispose of them. The purchasing power of an asset is ultimately predicated on a common confidence in the value of that asset; the danger for an investor is a reduction in or the disappearance of that confidence, as many participants in the dotcom bubble discovered. We leave it as an exercise for the reader to decide if the market price of diamonds, for example, is justified or whether, as an asset class, precious stones may be compared to Dutch tulip bulbs in the 1600s, a discussion of which may be found in Mackay (2003).

In the remainder of this book we shall see how an investor can use the time value of money and statistics to value a financial portfolio, understand the various risks to which it is exposed, and hedge those risks. Before we focus on valuation, and specific risks such as equity or credit risk, however, it is useful to spend some time considering the economic environment in which companies operate, the role of interest rates and the money supply, the nature of the business cycle, and the operations of central bankers such as the Federal Reserve. Much as a sailor needs to understand weather patterns in order to safely navigate the seas, an investor needs to be aware of the forces at play in the broad economy that can impact the price—and value—of his investments.

In this chapter we make no pretense of summarizing the entire canon of economic knowledge; rather, we highlight those aspects with which a market

¹To avoid tortuous linguistics, the authors use the masculine pronoun to refer to both males and females.

participant should have at least a passing familiarity. For a more detailed exploration, we refer the reader to any standard text in what Thomas Carlyle termed the dismal science.

Supply and Demand

The principles of supply and demand are key to understanding the behavior of producers, consumers, and investors. The *law of demand* states that, *ceteris paribus*, consumers will demand more of a product as the price drops. The *law of supply* states that more of a good will be produced as the price rises, again assuming all other things being equal. When consumers require more of a given product, its price will tend to increase, reflecting scarcity; producers will then have an incentive to increase production to satisfy the increased demand, tending to bring the price down. Conversely, if consumers require less of a product, the price will generally fall, producers will produce less, and ultimately the price will increase, reflecting the reduced supply. When the market is in equilibrium, the amount required by consumers is the amount being supplied by producers, and prices are stable at a level that can clear the market.

The laws of supply and demand apply to all products, including money. In order to borrow money, the borrower must pay a price, in the form of interest. A bank charges interest on loans it makes and pays interest to depositors who in turn lend it money. If a bank requires more funds to lend out, it must increase the price it pays for borrowing by, for example, increasing the interest rate it pays to its depositors.

An enterprise that needs to borrow money can do so by selling bonds to the investing public. An owner of cash has a choice between consuming the cash today (i.e., purchasing goods and services) and deferring his consumption of the cash by lending it to the issuer of a bond by purchasing the security. The price paid to the bond purchaser (i.e., investor) will be in the form of the interest he receives. The loan will be repaid at maturity when the investor receives the par amount of the bond (assuming the issuer has not defaulted). If the investor had previously disposed of the bond, the ultimate purchaser stands to receive the par amount. In our example, the investor is a lender of cash, the issuer of the bond is the borrower, and the interest rate charged on the bond is, in effect, the price of current consumption.

If there is an excess of borrowers over lenders, then the market is not at equilibrium. To reach that point, the price of borrowing (i.e., the interest rate) must increase. When the interest rate is sufficiently high, more investors will be prepared to lend to companies by purchasing their bonds. As we shall discuss in Chapter 2, bond prices generally move inversely with interest rates. When issuers increase the interest rate on their bonds, they are in effect

lowering the price (in present value terms), which by the law of demand will have the effect of increasing the number of bond buyers (i.e., cash lenders).

Interest rates in the market tend to reflect equilibrium pricing, incorporating participants' views of credit, liquidity, inflation, and other risks. As the consensus expectations for these risks change in the market, so too will the price that investors require to lend money.

Money Supply

The *quantity theory of money* suggests that there is a relationship between the money supply in an economy and the prices of goods and services sold throughout the year. Economist Irving Fisher expressed the relationship in the form:

$$MV = PQ \quad (1.1)$$

where M is the Money Supply, V the Velocity of Money, P is Price, and Q is Quantity of Goods produced in the economy. Since V and Q are relatively stable, a change in M will lead to a change in P ; in other words, an increase in money supply means that consumers will seek to purchase more goods and services, increasing demand and thereby prices—or, more succinctly, printing money usually leads to *inflation* with concomitant dangers for investors.

Inflation is a continual increase in the general level of prices of goods and services in the economy, leading to a reduction in the quantity of those goods and services that each unit of currency purchases. When lending money, therefore, for example, by purchasing a bond, an investor is compensated for anticipated inflation through the interest rate charged. Interest rates have two components, *nominal* and *real*. The nominal rate is merely the stated rate, for example, the bond coupon. The real rate is the yield effectively available to the investor after inflation:

$$\text{Nominal Rate} - \text{Inflation} \approx \text{Real Interest Rate} \quad (1.2)$$

If the nominal rate is 5%, for example, and inflation is 2%, then the lender's effective return will be approximately 3%. If inflation ultimately transpires to be higher than expected, then the lender will not have charged a sufficiently high coupon. If inflation is lower than anticipated, the borrower will have overpaid for the loan, all other things being equal.

An increase in inflation is bad for equity and fixed income markets, since the present value of future earnings and coupons, respectively, will be eroded, leading investors to require a larger inflation premium. Furthermore,

increases in expected inflation may lead to higher interest rates, resulting in increased borrowing costs for companies and individuals and a constriction of the economy.

Conversely, a decrease in inflation will generally be bullish for equities and fixed income instruments. If, however, the rate of inflation falls below zero, it can degenerate into *deflation*, a sustained decrease in the general price level of goods and services. Effectively a negative inflation rate, the result is an increase in the purchasing power of money. Deflation can be very harmful to an economy; in a deflationary environment, demand contracts because rational investors defer purchases, since their money's purchasing power increases over time, thus shrinking the economy and reinforcing the deflationary trend in a vicious circle of declining demand, increasing unemployment and stagnating markets. Debt burdens effectively rise, since notional debt principal remains unchanged while wages and profits fall, leading borrowers to curtail spending in order to be able to service their debts, further weakening the economy.

1.2 INFLATION AND UNEMPLOYMENT

In his eponymous curve, the economist William Phillips demonstrated the relationship between inflation and unemployment. As unemployment decreases, inflation tends to increase, and vice versa, the result of labor costs being a function of supply and demand; if labor is in short supply, wages and salaries will increase, providing an incentive to work. Conversely—in theory—if there is an oversupply of labor, then wages and salaries will be reduced, resulting in increased employment, as employers have an incentive to hire workers at cheaper wages. Unemployed workers tend to reduce aggregate demand in the economy, with the natural response being a reduction in prices, leading over time to increased demand, requiring producers to hire more workers.

In equilibrium, there should be no involuntary unemployment.² However, government intervention results in market inefficiencies. If there is a statutory minimum wage level (i.e., a price floor for labor), then supply and demand will not be equal and there will be no clearing price for wages, resulting in a surplus of labor (i.e., unemployment).

Figure 1.1 shows the relationship between the *consumer price index* (CPI), that is, the inflation rate, and the unemployment rate. In times when

²In an economy with relatively high levels of unemployment benefits, some workers may choose to cease working rather than accept lower wages.




FIGURE 1.1 As inflation—and interest rates—rise, unemployment tends to increase with a lag of two to three years. This trend, however, appears to have broken in early 2008.

the inflation rate is extremely volatile, such as 1970s and 1980s, the relationship is evident. In the late 1990s and early 2000s, with less volatility in the inflation rate, the relationship is less clear. It should be noted that extreme economic shocks, such as those of 2007–2010, can lead to significant increases in unemployment without a contemporaneous increase in inflation. However, significant increases in unemployment can in their turn lead to spikes in inflation caused by governmental fiscal responses to economic crises.

1.3 CENTRAL BANKS AND THE MONEY SUPPLY

Monetary policy, which governs the supply and cost of money, is typically dictated in an economy by a nominally independent or quasi-autonomous *central bank*, which usually has a mandate to control inflation by setting interest rates, conducting open market operations, determining depository institutions' reserve requirements, setting margin requirements, and by other

means. The central bank usually also has supervisory powers over the banking sector, and often is the *lender of last resort* in an economy.

The *Federal Reserve System* (the Fed) is the central bank of the United States. Unusually, it has a second, conflicting, task of controlling unemployment. The Fed was created by the Federal Reserve Act of December 23, 1913, largely as a response to the panic of 1907 and a series of earlier panics and bank runs. The Act, amended in 1977, articulates the Fed's mandate

to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates.

The Federal Reserve consists of 12 regional banks and is controlled by a seven-member Board of Governors appointed by the U.S. president. The Federal Open Market Committee (FOMC), consisting of the seven governors and five of the 12 regional bank presidents, oversees the open market operations by which the Fed controls monetary policy.

Central banks such as the Fed can conduct monetary policy (i.e., decrease or increase the money supply to stabilize economic growth and reduce the risk of inflation or deflation) by changing the discount rate, buying or selling securities, or by changing banks' reserve requirements, which can dramatically alter the money supply through the fractional reserve nature of bank deposits and lending.

The Fed is the lender of last resort to depository institutions, mostly commercial banks, who can borrow at the "discount window" at the so-called *discount rate*, the only rate set directly by the Federal Reserve. By choosing to raise or lower the discount rate, the Fed can attempt to discourage or stimulate bank borrowing. This in turn affects consumer interest rates and credit availability and, therefore, consumer spending.

In practice, depository institutions avoid the discount window, the use of which can be interpreted as a sign of financial weakness. The common method for banks to meet short-term, especially overnight, liquidity needs is to borrow each other's balances at the Fed ("federal funds"). The rate at which this borrowing is done, the *federal funds rate*, is the shortest term and therefore, usually, the most volatile rate in the U.S. interest rate markets.

Central banks usually change reserve requirements and the discount rate infrequently. On a day-to-day basis, they effect policy through *open market operations*, buying and selling government and other financial securities in the open market, or by lending and borrowing collateral.

The Fed strives to keep the federal funds rate close to the *federal funds target rate* set by the FOMC, typically 25 to 100 basis points below the discount rate. While normally the target is a specific number, on December 16, 2008, the FOMC set a *range* of 0 to 25 basis points (the discount rate

was set at the same meeting at 50 basis points). The federal funds rate subsequently stayed close to 15 basis points.

These open market operations typically involve *repurchase agreements* (repos) and *reverse repurchase agreements* (reverse repos). A repo is in effect a short-term loan extended to a *primary dealer* by the Fed in exchange for a collateral of Treasury securities. The borrower promises to repurchase the collateral at the end of the loan, typically one to seven days later, at a higher price. Repos inject liquidity into the economy, and the FOMC uses them to lower the federal funds rate. A reverse repo, in which the Fed borrows from a primary dealer, drains liquid funds and thus increases the federal funds rate.

Investors loathe uncertainty; if the future level of inflation cannot be estimated with a high degree of confidence, market interest rates will be high to compensate investors for this uncertainty; the resultant high borrowing costs will tend to diminish economic activity. The Federal Reserve, therefore, strives to balance the unemployment rate against inflation and ensure price stability. If the Fed is concerned about the current level of inflation, it may seek to raise interest rates by selling securities at reduced prices. Investors attracted by the lowered prices—and higher yields—will purchase those bonds for cash, which the Fed effectively withdraws from circulation. As there is now less cash in public hands, its price—the interest rate—rises. Conversely, to reduce interest rates, the Fed purchases securities. To attract sellers, the Fed will increase prices; investors will receive cash in return for the securities, increasing the money supply and thereby reducing the price of cash—the interest rate. In addition to its open market operations, the FOMC meets at six-weekly intervals—and convenes in emergencies, if required, by conference call—to consider and, if necessary, adjust the target and discount rates.

The market response to the actions by the Fed may not always be immediately intuitive. When the Fed tightens, for example, short-term rates rise. However, investors may believe that the Fed action will reduce the risk of future inflation; investors may as a result move into bonds, thereby prompting a drop in long-term rates. Similarly, a decrease in rates in response to an economic slowdown often leads to a rally in equities, because the benefit afforded by reduced borrowing costs may more than offset the impact of a slowing economy.

Figure 1.2 illustrates the recent history of inflation and the federal funds rate. Many large financial institutions employ “Fed watchers” who devote a great deal of time to monitoring inflation and interpreting other economic indicators in an effort to divine the future actions of the Fed and profit accordingly. For example, they may use the *Taylor Rule*, which makes use of current and expected gross domestic product (GDP) and current and desired inflation rates to determine a target interest rate, as the basis of their forecast of the federal funds target rate.




FIGURE 1.2 The federal funds rate is highly correlated with inflation. Monitoring current inflation may help investors forecast the federal funds rate.

1.4 THE BUSINESS CYCLE

Economic activity is measured by *gross domestic product* (GDP), a measure of the goods and services produced domestically in the economy, regardless of the ownership of the resources used to produce them.³ The *output gap* is the difference between current GDP and GDP based on a long-term trendline. When current GDP is lower than trendline GDP, the economy and inflationary pressures have slowed. Conversely, when current GDP is higher than trendline GDP, inflationary pressures have increased, and there is a risk that interest rates will rise.

Over the long term, the level of business activity in an economy can be seen to fluctuate through a series of peaks and troughs known as the *business cycle*. The business cycle has five phases, beginning with an initial recovery, progressing through an early expansion and late expansion, into a slowdown and finally into recession.

³*Gross national product* (GNP) is a measure of goods produced by the permanent residents of a country, irrespective of where they are produced (i.e., domestically or internationally).

Interest rates tend to rise during an early expansion phase, reducing in a recession phase, while inflation has a tendency to rise with a late expansion, peak in a recession, and fall during an initial recovery. Equity prices tend to be *leading indicators*, signaling future economic conditions, rising with an initial recovery and falling with slowdown and recession cycles in response to changes in the level of inflation and interest rates. Meanwhile, as discussed above, an increase in inflation reduces the appeal of bonds as an investment. In an expansion phase, therefore, as inflation increases, bond prices tend to fall. They will rise, however, during a recession when interest rates fall. Understanding where the economy is in the business cycle will help investors identify inflection points and to gauge the likelihood of particular Fed actions. While it can be difficult to identify precisely where the economy is in a business cycle or how long a particular phase will last, an understanding of the relevance of economic statistics can help investors anticipate economic swings. An in-depth discussion of the business cycle can be found at the website of the Economic Cycle Research Institute (www.businesscycle.com).

While inflation erodes the value of financial assets such as stocks and bonds, other investment opportunities, such as real estate, and other physical assets, such as gold or commodities, tend to perform better in inflationary times, with raw materials performing better than finished goods.

It is clear from Figure 1.3 that there is a strong positive relationship between the inflation rate and the spot price of gold. Gold is traditionally




FIGURE 1.3 Historically, investors have used gold as a hedge against inflation.

considered a safe haven in times of economic uncertainty and inflation. When inflation was high, in the late 1970s, for example, gold prices rose dramatically. Conversely, in the late 1990s when inflation was muted, the investment returns of gold were modest. Other nonperishable commodities, including oil and gas, are also positively correlated with inflation; those that cannot be stored, such as agricultural commodities, can have a negative correlation with inflation.

By now it should be clear that it is important for investors to understand where we are in the current business cycle and what growth rates have historically prevailed at similar stages in the past, and they should adjust their investment and hedging strategies accordingly.

1.5 PREDICTING THE FUTURE?

Some market observers claim that through the use of *technical analysis* they can identify patterns in the historical movements of securities prices from which they can draw inferences about future price levels. In a rational world, this suggestion might seem absurd, with many academics professing a belief in *efficient markets*. There is, however, a growing body of research into an area called *behavioral finance*, which suggests that investors do not always behave rationally.

The *law of large numbers* is often used to extrapolate from the behavior of a sample to estimate the behavior of an entire population in disciplines such as electoral opinion polling, loan portfolio loss projections, and insurance mortality calculations.

What the law of large numbers suggests is that while the instantaneous behavior of individuals may not be possible to predict, it is possible to make considered judgments about the future behavior of the broad market by examining the past and identifying relationships between, for example, consumer spending and the unemployment rate.

Economic forecasting is an art as well as a science, combining statistics and psychology. Although present circumstances may resemble those prevailing in some previous time period, we must remember that human behavior is always subject to change, reacting in turn to macroeconomic, political, and technological change.

1.6 ECONOMIC INDICATORS

Economic forecasting was recognized by Benjamin Graham in his classic 1940 book as providing “essential underpinning for stock and bond market,

industry and company projections” (see Graham and Dodd, 2002). The Fed uses selected statistics about the current state of the economy to set monetary policy. At the same time, many companies make use of these statistics, referred to as economic indicators, to produce forecasts of future economic conditions to guide their business strategy. Financial analysts also use forecasts to project future cashflows and interest rates, and asset managers use them in determining their asset allocations. Indeed, to some extent, an expectation of a given set of economic conditions becomes a self-fulfilling prophecy: If the broad market expects a favorable economic climate, for example, it will increase the likelihood of a benign environment. Conversely, as we have seen in 2007–2010, fear and uncertainty in the bond and equity markets can deepen a crisis.

So much attention focused on the release of key important economic reports can lead to significant swings in the market due to the interpretation of those data by market participants who use it to make inferences about the prevailing level of unemployment, inflation, the current state of the business cycle, and future actions to be taken by the Fed.

It is therefore incumbent on market participants responsible for managing and hedging investment portfolios to monitor key indicators as they are released and to understand the significance of leading, coincident, and lagging indicators. Leading indicators tend to move in advance of the general level of economic activity, coincident indicators illustrate the current state of the economy, and lagging indicators tend to follow behind the broad economy, confirming trends.

As with all data, the user ought to be aware of the time lag between the collection and distribution of the data, changes both in the composition of and the methodology used to calculate indices, and data measurement errors including survivorship bias and smoothing. A comprehensive discussion of various economic indicators can be found in Rogers (1994) and Yamarone (2004). A release schedule can be found at the U.S. Government Census Bureau website (www.census.gov/epcd/econ/www/indijun.htm). Another resource useful in gauging the state of the economy is the *Economic Report of the President*, a report prepared annually by the Chairman of the Council of Economic Advisors that provides a useful discussion of the economy tempered by pragmatic political realities. This report is available through the U.S. Government Printing Office (www.gpoaccess.gov/eop/).

The key indicator that investors focus on is GDP, which indicates the relative strength of economic activity. Sustained economic growth, evidenced through a steady rise in GDP, is welcomed by equity investors. Fixed income investors, meanwhile, examine the GDP number and related

information—such as the *GDP deflator*, a measure of inflation that can be computed by comparing real and nominal GDP—for signs of inflation or economic slowdown, often responding favorably to the latter. Consequently, it is not unusual for the bond markets and the equity markets to respond differently to the release of a given indicator.

In recent years, approximately 70% of U.S. GDP is accounted for by personal consumption, with about 10% of that coming from purchases of durable goods (those with a shelf life in excess of three years), 30% non-durable goods, and approximately 60% related to purchases of services such as housing, medical expenses, education, and recreation. Since personal expenditure is such a large component of GDP, investors seek to understand consumer behavior by paying close attention to indicators such as *Consumer Confidence*, which are scrutinized in an attempt to identify the likely trend of GDP. Other metrics of consumer behavior such as the level of *Car Sales* are measures of consumer optimism, as rational consumers ought to be reluctant to make expensive purchases such as of automobiles—often bought on credit—if they believe the economy is headed for a significant economic slowdown. Car sales can therefore prove to be a leading indicator of macroeconomic performance.

Durable Goods Orders similarly reflect broad economic sentiment, manifested by the willingness of both consumers and companies to purchase appliances and machinery. For example, a decrease in durable goods orders can often signal an imminent economic downturn.

The *ISM Manufacturing Index* provides a barometer of manufacturing activity. Like metrics of consumer behavior, this indicator is available on a monthly basis, in advance of the GDP number that is published quarterly, and so can be used by analysts to provide an early estimate of the likely level of GDP.

Pending Home Sales is another useful leading indicator, as house sales reflect the availability of credit and positive consumer sentiment and typically result in significant additional consumer spending as purchasers furnish their new homes.

The *Employment Situation Report* provides details of unemployment and wage inflation, and so is examined in an attempt to predict likely Fed actions. Other information provided in this report, such as the average number of hours worked, can serve as a leading indicator of future manufacturing activity and employment levels. The *ADP Employment Report* is available a few days in advance of the Employment Situation Report and can be used to estimate the level of the latter.

The *EIA Petroleum Status Report* reveals the supply and demand for oil, suggesting the direction of pressure on oil and energy prices, with a

concomitant influence on business profitability and providing insight into the underlying strength of the economy.

Market dislocations often happen when information is released unexpectedly, or, when released, is at variance with expectations. When economic statistics are published, the importance is often not in the number itself, but rather in how that number deviates from what the market had been anticipating and had previously priced in. Similarly, revisions to statistics previously released can move the market, particularly if indicative of a trend. If employment numbers come in worse than had been expected, that can be a bullish signal for bonds, lowering yields, on anticipation that the economy is weaker than previously assumed and an expectation that the Fed may stay or even lower rates. Figure 1.4 illustrates the market consensus expectations at two dates in July 2009 for the Fed Funds target rate to be set by the FOMC at its next meeting.

Managers of financial portfolios should try to anticipate market sentiment and reactions, and position their portfolios accordingly. They should also examine the political landscape and analyze the actions of the Administration. For example, in August 2009 the Congressional Budget Office estimated that the budget deficit for 2009 would exceed 11%. With federal programs such as Social Security and Medicare having unfunded liabilities in excess of \$100 trillion, the risk of a significant increase in interest rates and inflation is substantial (Villarreal, 2009). Portfolio managers should analyze




FIGURE 1.4 Over time, the market's expectation of the federal fund rate can change.

Source: Bloomberg Finance LP

TABLE 1.1 Key U.S. Economic Indicators and Their Reporting Dates

Economic Report	Approximate Release Date
Car Sales	1st or 2nd business day for prior month
ISM Manufacturing Index	1st business day for prior month
Pending Existing Home Sales	2nd business day for two months prior
Employment Situation	1st or 2nd Friday for prior month
ADP Employment Report	2 days prior to Employment Situation report
EIA Petroleum Report	Wednesday for prior week
International Trade	9th-16th for two months prior
Retail Sales	11-14th for prior month
Housing Starts	16-20th for prior month
Producer Price Index	2nd or 3rd week for prior month
Industrial Production/Capacity Utilization	Mid-month for prior month
Consumer Price Index (CPI)	2nd or 3rd week for prior month
Durable Goods Orders	22nd-28th for prior month
Consumer Confidence	Last Tuesday of the month
Personal Income/Consumer Spending	22nd-31st for prior month
GDP	Last week for previous quarter

the potential risks and benefits to both their home and the global economies by governmental actions and choose their investment and hedging strategies appropriately.

Some commonly used economic indicators and their approximate release dates are listed in Table 1.1.

CHAPTER 2

Risk: An Introduction

*Tzu-hsia said, “Learn widely and be steadfast in your purpose,
enquire earnestly and reflect on what is at hand, and there is no
need for you to look for benevolence elsewhere.”*

Confucius, The Analects

2.1 WHAT IS RISK?

We can define risk in many ways, including as the possibility of uncertain change, particularly a change for the worse. From a financial perspective that means that there is a chance of suffering an economic loss. When investors talk of individual security risk, they may be thinking of the risk of a corporate bond defaulting, or of a stock price declining. When referring to portfolios, risk is usually taken to mean that there is a chance of some return objective not being met. That objective can be an absolute return target, such as an annual return of 10%, or it could mean “not losing money,” that is, maintaining the current dollar value of the portfolio. Alternatively, for a pension fund with defined future liabilities to be met, it could be a relative return requirement—for instance, that the portfolio assets match the portfolio liabilities over time. With an absolute return requirement, the return objective is known, with certainty, in advance. With a relative return requirement, such as meeting or exceeding a benchmark return, the *constraint* on performance is known in advance with certitude, but the return objective is unknown.

Some risks can be determined quantitatively, as we shall explore in later chapters. For example, when purchasing a ticket in a typical lottery, the odds of winning can be determined with absolute precision to be, for example, one in one million, although the payout may be uncertain. If the lottery jackpot is guaranteed to be exactly \$1 million after taxes, then buying a

lottery ticket for a price of \$1 is an indifferent proposition, but only if one can be sure that the jackpot will not be shared. From a rational perspective, a player should be indifferent to the proposition, because he could buy all the tickets at a cost of \$1 million and be sure of being returned \$1 million. Any larger *guaranteed* payout to him makes this an attractive investment. Were one to purchase only 999,000 tickets, the risk of an economic loss would be 1 in 1,000 or 0.1%.

Other risks can only be *estimated*. For example, a person's risk of being struck by lightning cannot be known with certainty. If we use our observations of weather and other factors such as the number of lightning strikes per year, geographic area, and population density, we can make an estimate predicated on *assumptions*. In this case, we have used past history to make projections about the future. Similarly, by examining the historical returns of two stocks, A and B, we might find that both stocks had mean annual returns of 12%. Both stocks might appear to be reasonable investments but this judgment is based only on the *average* return. If we look at the *standard deviation* of returns, we might see that A had a standard deviation of 5%, and B a standard deviation of 20%. Standard deviation is a measure of how a series, in this case returns, fluctuates. As investors, when faced with two investment alternatives with the same mean return, we should rationally opt for the one with lower standard deviation to minimize the risk of negative surprises.

There is a serious danger that should not be overlooked when examining returns and standard deviations. Future returns cannot be known *ex ante*; they may only be estimated. The prudent manager must be careful not to assume that conditions that prevailed in the past will continue to do so in the future. Just as a person's risk of being struck by lightning can increase as the incidence of thunderstorms increases, the variability of returns in the future can increase. We must be careful to remember that history includes only all those events that have happened, not necessarily all of those events that *could* happen. This is particularly important when attempting to calculate risk exposures, for example, using a technique such as *Value at Risk* (VaR). We shall explore these topics further in Chapter 5.

It behooves managers to focus not merely on a quantitative analysis of risk, but also on a qualitative analysis. They should focus not just on what has happened, but also on what can possibly happen in the future. In the words of Thoreau, "To know that we know what we know, and that we do not know what we do not know, that is true knowledge."

Every asset is exposed to market risk, meaning there is a possibility that its value will decrease due to a change in the value of one or more market variables, including interest rates, foreign exchange rates, equity prices, and commodity prices. For some assets, the market risk is clear. The price of a

barrel of oil, for example, fluctuates continually due largely to supply and demand driven not just by industrial and consumer demand but also the impact of speculators.

For other assets, such as corporate bonds, the price can vary due not only to supply and demand, but also to changes in interest rates. In addition, the price can move due to factors such as a deterioration in the perceived credit quality of the issuer, perhaps the result of a change in opinion from a rating agency. Prices may also move due to less obvious factors such as a reduction in available market liquidity, or changes in market volatility.

In this book we shall concern ourselves primarily with the risk of financial assets, those whose value derives from a contractual claim, such as stocks, bonds, bank deposits, currency, and financial derivative instruments, but we shall from time to time refer to non-financial assets, such as real estate and commodities.

2.2 RISKS OF FINANCIAL INSTRUMENTS

While a manager is ultimately concerned with the risks of his portfolio as a whole, he needs to understand the dangers associated with the underlying portfolio constituents. In the end, it is the risks of these individual securities and, especially, the *interplay* of these risks, that shape the portfolio's risk profile.

Investors in securities issued by a company are faced with both individual company risk, referred to as *non-systematic risk*, or *specific risk*, and the risk associated with the market as a whole, known as *systematic risk*. In the remainder of this chapter we shall highlight many of the risks, some obvious and some less obvious, that can directly impact the value of a portfolio including indirect risks, such as counterparty risk, legislative risk, et cetera.

Some of those risks can readily be hedged, directly or indirectly, while others can only be managed, and not hedged.

Equity Risk

A rational investor who purchases stock in a company is seeking an investment return, manifested as some combination of price increase—*capital gain*—and *dividend income*. For there to be a capital gain, there must at some future time be a purchaser willing to pay an increased price.

The price an investor chooses to pay may be based on a quantitative analysis using a model such as a dividend or earnings growth model. If changes in the fortunes of the company lead to lower dividends or diminished growth prospects, the price that a subsequent purchaser is prepared to offer

may be lower, resulting in a loss for the investor. In an extreme case, the company may file for bankruptcy protection or be nationalized, *de jure* or *de facto*, with the stock often ending up worthless.

Alternatively, rather than being based on a quantitative analysis, the price the investor is willing to pay may be the result of pure speculation, based on a belief, often naïve, in the existence of a *greater fool* who will in the future be prepared to pay a higher price for the stock. This type of speculative investing is common in market bubbles; alas, from the speculator's perspective, bubbles invariably burst, with an often dramatic reduction in prices. For a classic exposition of this subject, the reader is referred once again to Mackay (2003).

Interest Rate Risk

The prices of most fixed income instruments are negatively correlated with interest rates. When rates rise, for example, the price of most fixed rate bonds¹ will decrease since the yield that the market now requires makes them less attractive investments. Similarly, a widening of prevailing corporate credit spreads will lead to a decrease in corporate bond prices due to the effective rise in bond yields. In both cases, investors required to subsequently sell a bond prior to its final maturity will suffer a capital loss. A fall in market rates, however, or a tightening of credit spreads, will tend to increase the price of bonds. The longer the maturity of the bond, the greater the price sensitivity to changes in interest rates, since the cashflows are discounted for more periods. The higher the bond coupon, however, the lower the price sensitivity to changes in interest rates, since the principal due at maturity is a smaller proportion of total bond cashflows than for a lower coupon bond.

A convenient measure of the sensitivity of bond prices to changes in interest rates is duration, expressed in years. The *modified duration* $\mathcal{D}$ of a bond—generally referred to nowadays as “duration,” in place of an older measure known as “Macaulay duration”—is the sum of the present value of all the bond cashflows, weighted by the length of time to their receipt, divided by the present value of cashflows, divided in turn by unity plus the bond's yield:

$$\mathcal{D} = \frac{1}{P} \sum_{t=1}^T t \frac{CF_t}{(1+y)^{t+1}} \quad (2.1)$$

¹A notable exception includes *interest-only* bonds, which we shall consider in our treatment of structure risk later in this section.

where CF_t is the cashflow at time t , y is the yield to maturity, P is the price of the bond, and annual compounding is assumed.

For small interest rate changes (e.g., 1%), the change in price of a bond is equal to the negative of the change in yield to maturity of the bond times the modified duration, so that modified duration is a reasonable measure of the interest rate risk of a bond. As an example, if the modified duration of a bond is five years, then if rates increase 1% the change in price would be:

$$\frac{\Delta P}{P} = -\mathcal{D} \times \Delta y$$

$$\text{Fractional Price Change} = -5 \times 1\% = -5\% \quad (2.2)$$

Modified Duration is a percentage price change for a 1% change in interest rates; to determine the dollar impact of a one basis point decline in rates, we derive the *dollar value of a basis point (DV01)* as follows:

$$DV01 = \mathcal{D} \times P \times 0.0001 \quad (2.3)$$

With certain kinds of bonds, such as mortgage-backed securities, the total cashflows that an investor receives can vary according to the level of interest rates due, for example, to changing prepayment behavior. For these bonds, *effective duration* is used when calculating price sensitivity to interest rate changes to account for the fact that changes in yield will change the expected bond cashflows. Another measure of duration is the *spread duration* of a bond, which measures the sensitivity of its price to a change in the bond's spread over its benchmark.

The coupon on bonds that have a floating rate adjusts in response to changes in interest rates. Their price sensitivity to interest rate changes is due to the fact that there can be a delay between the change in market yields and the next reset date of the bond's coupon. The modified duration of such a bond is defined to be equal to the interval between reset dates.

As discussed above, duration tends to increase as the coupon decreases. For a *zero-coupon bond*, in fact, the modified duration is equal to the remaining maturity of the bond. It is worth noting in passing that the price of a particular kind of mortgage-backed security known as an *interest-only bond (IO)*, one that receives only interest and no principal, typically increases as interest rates rise, and so has a *negative* duration.

The modified duration of a bond is a good measure of bond price sensitivity to interest rate changes for small changes in yield. However, for large changes in yield the price decline from a yield increase is usually less than the price increase for a yield reduction of the same magnitude.

This phenomenon is due to the *positive convexity* of the bond. Callable bonds, which can be called by the issuer, exhibit negative convexity at low interest rates. For investors, positive convexity is generally an attractive bond feature, with negative convexity being undesirable. This topic is discussed in more detail in the discussion of call risk later in this section.

Convexity C can be incorporated into our previous formula to produce a better estimate of a bond's price change for a given change in yield:

$$\frac{\Delta P}{P} = -D \times \Delta y + \frac{1}{2} \times C \times (\Delta y)^2 \quad (2.4)$$

The reader can find a detailed discussion of these and other fixed income risk measures in, for instance, Tuckman (2002).

It is not just fixed income instruments that are exposed to interest rate risk. For small interest changes, equity prices can move in tandem with Treasury yields as investors view modest rate hikes as a sign of a robust economy. As market yields continue to increase, however, or credit spreads widen, companies will have less incentive to make capital investments due to the additional interest that they would have to pay on new debt, thereby reducing potential growth in their future earnings. In addition, their future earnings and any dividends they pay must now be discounted at a higher rate, reducing their present value, with a consequent reduction in the equity price a rational investor would be prepared to pay.

Other asset types, such as real estate and commodities, are also sensitive to interest rate moves.

Yield Curve Risk

The modified duration of a bond or portfolio measures its price sensitivity to a parallel yield change across the curve. However, fixed income investors are also subject to risk that the shape of the yield curve will alter in a non-parallel fashion, for example, flatten, become steeper, or even invert. This is particularly important when an investor's portfolio of assets and liabilities exhibits a duration mismatch.

A bond's *key rate duration* measures its price sensitivity to a change in the rate at any of a series of key points, typically 11, along the yield curve. This allows an investor to compare, for example, the impact on both his assets and his liabilities of a shift in the 10-year rate, with the rest of the yield curve unchanged.

Consider two portfolios, A and B. Portfolio A consists of \$100 in a 10-year zero coupon bond. Portfolio B consists of \$50 in cash and \$50 in a 20-year zero. Since the duration of a zero coupon bond is equal to

its maturity, and since cash (or a cash-like instrument), being insensitive to interest rates, has a duration of functionally zero, then the modified durations are:

$$\text{Portfolio A: } (0/100) \times 0 + (100/100) \times 10 + (0/100) \times 20 = 10$$

$$\text{Portfolio B: } (50/100) \times 0 + (0/100) \times 10 + (50/100) \times 20 = 10$$

For both portfolios, the modified durations are the same. However, note that Portfolio A is sensitive only to a change in the 10-year rate; if that rate increases by 1%, then the portfolio will decrease in value by 10%, whereas Portfolio B, with no exposure to the 10-year rate, would be unaffected by the change.

If the interest rate sensitivity of the liabilities in a portfolio does not match those of the assets, a change in rates can cause liability obligations to exceed assets, causing for instance a previously appropriately funded pension plan to become underfunded. The *Savings and Loans Crisis* of the 1980s was due in large part to duration mismatching. Savings and Loan Associations (S&Ls) issued long-term mortgages to their customers financed by short-term deposits. As interest rates spiralled, the cost of the liabilities increased while the rate on the long-term mortgages remained unchanged, resulting in massive deficits. In 1980, for example, net S&L income totaled \$781 million. In 1981 it fell to −\$4.6 billion and to −\$4.1 billion in 1982, resulting in just three years in the failure of 118 S&Ls, 493 voluntary mergers, and 259 supervisory mergers, and ultimately an estimated cost of \$160 billion. A detailed treatment of this saga may be found in FDIC (1997).

Credit Risk

Investors in fixed income securities are exposed to the risk that the issuer will not be able to pay interest as it falls due and to repay principal by the maturity date of the instrument. The severity of the loss suffered by the investor, often referred to as *loss given default*, is a function of how much principal has been repaid before the default occurs, and how much, if any, recovery will be made in the event of a default. With *bullet bonds*, the entire principal amount is due to be repaid at maturity, so any default results in a loss, before recovery, of the entire principal amount.

With *amortizing bonds* and bonds with *sinking fund* provisions, principal is often retired throughout the life of the bond and so the severity of a loss may be reduced. Amortizing principal schedules are a common feature of *asset-backed securities* (ABS). It should be noted that with asset-backed securities, the credit risk of the bond is related to the performance of the

pool of collateral supporting the bond and not the credit quality of the issuer. Often the rating of an asset-backed bond can be higher than that of the issuer—at least at the time of bond issuance. Subsequent performance of the collateral, especially sub-prime mortgages, for example, may sometimes prove somewhat disappointing.

Convertible bonds combine characteristics of equity and fixed income securities. Holders of convertibles are more directly exposed to the credit risk of a firm than are equity owners, as they are entitled to the timely payment of interest and, in the event of the convertible being *busted*, ultimate payment of principal. Investors in convertibles and equity investors can both be considered to be exposed to increased credit risk if the share price underperforms; if the debt instrument does not convert to equity, the company will have to repay principal, reducing its financial flexibility.

Bonds with put, call, or sinking fund provisions introduce other nuances to the investment decision. Puttable bonds (e.g., Liquid Yield Option Notes (LYONs)) provide a benefit to the investor if interest rates rise, since he has the right to sell the bond back to the issuer at par despite any change in yield. Such an investor can, however, be exposed to a greater degree of credit risk than that associated with *plain vanilla* bonds, as this bond feature can place additional stress on the issuer, whose financial flexibility may be constrained if required to redeem bonds before their anticipated final maturity. Similarly, investors banking on their bonds being called, or redeemed via a sinking fund provision, may on occasion learn to their chagrin that the issuer is unable to execute the early redemption provision.

Investors in fixed-income securities are also exposed to the risk that the bond's *credit spread* (i.e., its required yield over its benchmark, which might be the swap curve or the on-the-run Treasury of the same maturity) might widen out, thereby reducing the price of the bond. This widening of the credit spread might be due to a perception in the marketplace that the credit quality of the bond has deteriorated, or it could be due to a widening of spreads for all bonds in that sector or at that credit rating, due to a *flight to quality* by investors in general away from lower credit issues into higher quality or government bonds.

Reinvestment Risk

An investor who purchases a zero-coupon bond knows with certainty the yield he will receive on his investment, barring a default. An investor who purchases a *coupon-bearing* bond, however, does not know the total return he will receive over the lifetime of the bond. This is because if interest rates drop before the maturity of the bond, the investor may not be able to reinvest the coupon payments at the stated coupon of the bond. As coupons increase,

reinvestment risk rises. It should be noted that “risk-free” coupon-bearing government bonds, therefore, are not risk-free investments.

Basis Risk

When an investor uses one instrument as the basis for hedging a second, non-identical, one, he is exposed to the risk that the price changes and risk profiles of the two instruments will not exactly offset each other. Consider an investor who attempts to hedge an equity portfolio using S&P 500 futures. He will face basis risk unless his portfolio consists of the same stocks, in the same proportions, as the S&P 500. Investors who attempt to hedge their portfolios may, because of a basis mismatch, eliminate some but not all of the underlying risks in their portfolio. For example, an investor who hedges the duration of a portfolio of corporate bonds using Treasury futures can reduce or eliminate interest rate risk by matching durations, but he is still exposed to spread risk. Similarly, an investor who hedges his position in a mutual fund through equity futures will still be exposed to redemption, manager, and other risks that we explore in the remainder of this section. In many cases basis risk does not create additional risk exposure, but rather relates to an imperfect hedge failing to eliminate all risk factors. In some cases, however, exposure to basis risk may lead to economic loss greater than for an unhedged position. For example, Figure 2.1 illustrates how in 2009 an investor who attempted to hedge a position with equity exposure to the financial sector using a triply levered financial sector bull market or bear market exchange-traded fund (ETF) would have seen disastrous results.

Basis risk can develop over time in portfolios originally perfectly hedged. This is especially true when hedging with levered instruments such as derivatives. This topic is explored further in our discussion of *delta hedging* in subsequent chapters.

Call Risk

An investor who purchases a callable bond is exposed to the risk that the issuer will take advantage of the call provision to redeem the bond early, effectively shortening its maturity. This will usually happen when interest rates have decreased, and so exposes the investor to reinvestment risk. *Mortgage-backed securities* (MBS) have implicit call provisions, due to the ability of the underlying borrowers to prepay their mortgage loans in response to favorable refinancing opportunities. A callable bond is worth less than a plain vanilla bond, all other things being equal. Conceptually, the reason is very clear—the issuer of a callable bond (or a mortgage borrower, in the case of an MBS) has retained an option that gives him the right to call the bond




FIGURE 2.1 Hedging a Russell 1000 Financial Index exposure using the Direxion Daily Financial Bull 3x Shares (FAS) or Daily Financial Bear 3x Shares (FAZ) can introduce basis risk.

(or prepay the mortgage) in favorable situations. This option has economic value, so the price of a callable bond should equal the price of a comparable plain vanilla bond *less* the price of the option.

Bond prices have an asymmetric response to changes in interest rates. For most non-callable bonds, a decrease in market yields will result in a price increase that is greater than the price reduction that would occur for an increase in market yields of the same magnitude. This phenomenon, as discussed above, is referred to as a bond's *positive convexity*.

For callable bonds, the increase in price due to a decrease in market yields is tempered by an increased likelihood that the bond will be called, or in the case of a mortgage-backed security, that the underlying mortgages will be prepaid. These bonds therefore exhibit *negative convexity* at low interest rates, with the price of a callable bond unlikely to trade significantly above par. Callable bonds' negative convexity at low interest rates has one beneficial consequence for the holder: An increase in yield from a low level to a new one still below or near the coupon would result in a smaller price drop than in the case of a plain vanilla bond. The penalty, of course, is a smaller gain if the rates decline.

Foreign Exchange Risk

A fixed-income investor who purchases a bond denominated in a foreign currency will be faced with the risk that his home currency will appreciate before the principal amount of the bond is redeemed. Similarly, equity investors who invest in shares of companies listed in exchanges outside their home countries (e.g., U.S. investors who buy *American Depositary Receipts* (ADRs) that represent ownership interests in foreign companies), or those who invest in companies the bulk of whose business is in other currencies, are faced with the risk of their home currency appreciating relative to the foreign currency.

Insurance Risk

Issuers of fixed-income securities, especially in the United States, sometimes take advantage of guarantees written by highly rated insurance companies to increase the credit rating of their bonds, to reduce the yield demanded by investors. A default by or downgrade of the insurance company providing the guarantee could result in a lowering of the rating of the bond, resulting in a higher required bond yield and lower price. In addition, the lowered rating might cause the bond to be in contravention of the purchaser's investment criteria, requiring the bond to be disposed of forthwith, at a loss.

If a bond insurer has insured a large portfolio of bonds, many of which have highly correlated risk exposures, then there is a risk of *contagion*. This risk was clearly demonstrated during the credit crisis of 2007–2010, when the poor performance of the collateral supporting a large number of insured asset-backed securities led to the failure of several monoline insurance companies that had provided guarantees to the ABS investors. With the failure of the guarantor, even uncorrelated securities such as municipal revenue bonds that, through the use of an insurance policy, had achieved a credit rating higher than that of the issuing municipality, suffered a reduction in credit rating.

Structure Risk

As mentioned previously, certain types of fixed-income instruments are structured so that the interest and principal payments are supported by cashflows generated by a discrete pool of collateral. The collateral underlying these *asset-backed securities* may consist of mortgages from which the cashflows *pass through* the MBS relatively undisturbed, or into a *collateralized mortgage obligation* (CMO) in which the cashflows are *tranche*d into various classes, each of which can have different coupon, average life, and

final maturity characteristics. Other forms of collateral found in asset-backed securities include home equity loans, auto loans, credit card receivables, movie box office receipts, intellectual property and royalty payments, and 12b-1 fees on mutual funds. In short, pretty much any stream of reasonably predictable future cash flow receipts can be structured into an asset-backed security.²

There are a number of characteristics of asset-backed bonds that an investor should consider. First, the credit-worthiness of the bond is predicated on the behavior of the underlying collateral. If the performance of a pool of mortgage loans, for example, is worse than anticipated (i.e., if there is a significantly higher than expected level of mortgagor default), then the ultimate performance of the bond will be worse than might reasonably be expected from the initial rating of the bond. Similarly, if the mortgage borrowers repay their loans in a pattern that differs significantly from the at-issue expectations for the bond, then the weighted average life, final maturity, average coupon, and total return may vary from what had been anticipated, a result of *prepayment risk*, in the case of shortening, or *extension risk*, in the case of lengthening, of the bond's final maturity. From a credit risk perspective, it is more desirable to have loans prepaid rather than have the loans extend due to slower than modeled repayments; a loan that has prepaid cannot subsequently default, whereas if economic conditions lead to a large number of loans extending, they may ultimately result in an increase in the number of those loans subsequently defaulting.

Second, in a tranching structure, the *cashflow waterfall* may allocate cash receipts from the collateral among the various classes in ways that may vary depending on the performance of the underlying collateral due to actions of *performance triggers* in the waterfall.

Third, *collateral servicing* is a critical element of the creditworthiness of the security. Any failure on the part of the asset servicer, who is responsible for collecting payments due on the underlying collateral and for monitoring the behavior of collateral obligors, may manifest itself in degraded deal performance. This can happen when banks or other servicing groups merge and rationalize their operations.

Fourth, any deficiencies in *due diligence* on the part of the issuer or the rating agencies, or errors in structuring, may lead to performance of the bond that is worse than initially anticipated.

Fifth, the type of any *credit enhancement* in the structure needs to be understood. As the adage suggests, cash is king. A dedicated, pre-funded,

²Whenever a model is used to project cashflows in a transaction, one must consider the exposure to *model risk*.

cash collateral account to cover anticipated losses is preferable to the use of *excess spread* or *over-collateralization*, since cash cannot default. As a collateral pool starts to deteriorate, for example, as borrowers start to become delinquent on a pool of auto loan receivables, excess spread will start to vanish rapidly. As those delinquencies turn into defaults, credit enhancement in the deal in the form of over-collateralization will start to reduce due not just to loans in the financed pool defaulting but as loans that formed part of the credit enhancement start to default as well. In cases where the default rates on the collateral are higher than anticipated at the time the deal was structured, the deal may effectively have substantially less credit enhancement than was originally assumed.

In certain kinds of collateralized mortgage obligations, the payment streams from the underlying mortgages may be disaggregated into their constituent interest and principal components. These distinct payment streams may then flow into segregated classes of the deal, referred to respectively as *interest-only* and *principal-only* classes (IOs and POs). The behavior of the IO and PO classes can vary significantly with prepayments. With lower than expected prepayments, for example, the IO class will receive more interest payments than originally modeled; the PO class, however, will have a longer final maturity than expected. Since prepayments typically decrease when interest rates rise, the value of IOs increase with increasing rates. This makes IOs one of the few investment instruments with *negative duration*.

Asset-backed bonds are also exposed to *concentration risk*. Similar to correlation risk, which we discuss below, this risk arises from the fact that the collateral supporting the bond may share similar demographic characteristics. For example, an ABS bond may be backed by mortgage collateral consisting of loans made to borrowers in a particular state. If there is a localized economic downturn, the performance of the bond may be disproportionately impacted compared to otherwise similar bonds with diversified collateral.

Investors should also be aware of the existence of a class of securities referred to as *structured notes*. Unlike asset-backed bonds, these securities are usually not backed by collateral, but instead are direct obligations of the issuer. The amount of interest and any principal that they may pay can be the result of complicated formulae, with coupon rates that vary directly—or even inversely—with the level of some benchmark rate such as an interest rate or an equity index. In addition, the coupon rate can change as a multiple of changes in the benchmark rate, or may move discontinuously with changes in the reference rate. Recent evolutions of this class of product, now commonly referred to as *structured products*, include those whose return is based on the dividend of one or more stocks, or the return of the best or worst performer of a basket of stocks, commodities, or currencies—in

short, the variety of product offerings seems limited only by the imagination of financial alchemists. While some of these relatively new instruments guarantee the return of most or all of the invested principal, it should be noted that some of these products incorporate features such as a *put spread* or a *knock-out option*, which offer only limited downside protection, or, even worse, protection that disappears entirely if the reference instrument or index decreases more than a specified amount. Consequently, it is imperative that an investor in such instruments have a thorough understanding of the payout formulae.

Model Risk and FAS 157

Many highly complex *structured products* once initially purchased are not frequently traded in the secondary market. In consequence, it is often not possible for a firm to obtain external prices to *mark to market* (MTM) the security on a periodic basis. As a result, the firm may use a computer model, often proprietary and non-transparent—a *black box*—to replicate the transaction and project the cashflows in order to generate for the deal a current *fair value*, that is, the price at which an asset could be exchanged, or a liability settled, between willing market participants, other than in a liquidation.

To ensure some degree of conformity in the approaches used to value securities the Financial Accounting Standards Board (FASB) in 2006 promulgated Statement of Financial Accounting Standards No. 157, “Fair Value Measurement.”

The standard distinguishes between those observable inputs obtained from sources independent of the reporting firm and those unobservable inputs, which by necessity are based on assumptions, used in determining fair value. Based on the nature of the inputs, there is a hierarchy of three levels:

- A **Level 1** asset or liability is one that is priced using an unadjusted quoted price for an identical asset or liability in an active market. Liquid exchange-traded stocks and most government debt issued by major developed economies fall into this category.
- A **Level 2** asset or liability’s value is either (a) a quoted price for a *similar* asset or liability in an active market; or (b) a quoted price for an identical or similar asset or liability in markets that are not active; or (c) an asset or liability whose value is derived from a model in which all significant inputs and value drivers are observable in active markets. Most derivatives on liquid stocks and indices as well as developed economy interest rate derivatives are Level 2 assets.

- **Level 3** assets and liabilities do not have quoted prices and are generally valued using a model; their values derive from valuation techniques in which one or more significant input or value drivers is unobservable. This level includes most structured products: MBSs, Collateralized Debt Obligations (CDOs), Collateralized Loan Obligations (CLOs), and other kinds of asset-backed securities; it also includes bank loans and generally anything that is not a liquid security or a “straightforward” derivative on such a security.

The Standard can be found at the FASB website (www.fasb.org/st/summary/stsum157.shtml).

Note that in periods of market dislocation, such as the credit crisis of 2007–2010, when market liquidity is scarce, there can be a significant disconnect between the extrinsic prices quoted in the marketplace and the intrinsic value that might seem reasonable, particularly to a hold-to-maturity investor. Model risk, therefore, can be accompanied by liquidity risk and market sentiment risk.

Model risk arises because a critical part of many models is the use of factors such as the movement of interest rates, equity or fixed income market volatility, and perhaps the behavior of individuals responsible for generating the cashflows associated with any collateral in the deal. Some of these inputs cannot be directly observed in a timely fashion and so must be estimated, that is, they are Level 3 inputs. There can be an incentive for a portfolio manager or trader to judiciously adjust the inputs to achieve an optimum valuation for the security. Furthermore, the model may contain inadvertent errors that result in the calculation of incorrect prices. Consequently, an investor should be aware that the price quoted for a security may not be robust.

Liquidity Risk

The ability to trade an asset at low cost, in a timely fashion, and with minimal market impact is dependent on the existence of sufficient liquidity in the market. Unfortunately, like over-collateralization, liquidity is ephemeral—it is there when not needed, but when it is required, it will, all too often, be found to have disappeared. This is particularly true of exotic asset-backed securities, structured notes, and, in certain circumstances like the liquidity crisis of 2007–2010, even plain vanilla corporate bonds.

As the recent liquidity crisis has demonstrated all too clearly, the disappearance of liquidity can become contagious *in extremis*. A severe widening of credit spreads for sub-prime mortgage-backed bonds and a growing realization of systemic problems in that market caused liquidity for those bonds

to dry up. Unfortunately, panic gripped the corporate bond and equity markets, fueled by lack of knowledge of counterparty exposure to troubled collateral and exacerbated by significant deleveraging by hedge funds in part prompted by margin calls, with that blind fear causing liquidity to disappear in most of the broad bond market.

Alas, lack of liquidity is a phenomenon that occurs in the market all too often. In 1998, for example, the Russian debt crisis triggered a liquidity trap that led to the collapse of the celebrated hedge fund, Long-Term Capital Management (Jorion, 2000; Lowenstein, 2000).

Lack of liquidity can cause dramatic price swings in the equity markets. If an investor borrows and sells shares that he does not own—a process known as a *short sale*—he will have to return the shares at some point in the future. If the price starts to rise in the meantime, he will have to meet a margin call and so will usually seek to repurchase the security in the open market. Unfortunately, other short sellers will be in the same predicament and so place further upward pressure on the stock price, the result of a *short squeeze*.

When a stock is thinly traded, which could be a consequence of a company's small size, or of a large proportion of the stock being closely held, relatively small trades can have a disproportionate impact on the stock price. Equity and fixed income portfolio managers who seek to add to or reduce their holdings of a security should carefully examine the size of their proposed trade in the context of the average daily trading volume and, if appropriate, execute their trade over a period of days to reduce the *market impact cost*.

Liquidity risk can be of concern to investors in mutual and hedge funds, particularly when they seek to exit the fund in times of scarce liquidity and when managers seek to balance the interests of those shareholders who are seeking to leave the fund against those who choose to remain. This risk, which we refer to as *redemption risk*, is linked to *manager risk* discussed below.

Managers of nonprofit and other funds such as *endowments* and *foundations* are subject to constraints on their investment strategies, and may be required to liquidate a portion of their assets each year. They may be obligated to produce current streams of income to meet current budgetary needs, and therefore must structure their portfolios so that there is sufficient liquidity to meet the particular requirements of their institutions.

Derivatives Risk

Derivatives, like chain saws, can be invaluable tools. Used improperly, they can also inflict severe injury, sometimes resulting in death. Derivative

instruments can be used for hedging purposes, such as protecting against the risk of a decline in the equity markets or an adverse change in interest rates. They can also be used for speculative purposes, both in relatively benign fashion such as purchasing call options, and in very dangerous fashion such as writing *naked calls*—selling a call on a stock you do not own.

Even more insidious than the conspicuous misuse of derivatives is a speculative application of derivatives that is concealed from the owners or trustees of the invested funds. For example, Robert Citron, as Treasurer of Orange County, California, entered into large interest rate derivative contracts with notional amounts significantly in excess of the value of the Orange County funds. While these contracts were written to benefit from a reduction in interest rates, when the Federal Reserve raised interest rates, the resultant losses forced Orange County to declare bankruptcy.

Used appropriately, derivatives can be used to reduce losses or to increase returns. For example, in a volatile equity market an investor may purchase protective puts to minimize downside exposure. In a relatively flat equity market, the writing of *covered calls*—selling a call on a stock you own—may be used to generate premium returns. Note, however, that purchasing downside protection is like purchasing most other kinds of insurance—there is an associated cost. Similarly, when writing covered calls, there is a potential *opportunity cost* to the strategy; if the market were to stage an extensive rally, the upside gains would be lost to the call writer through an exercise of the call—with the foregone profits being mitigated in part, of course, by the premium received.

While mutual funds make only modest, if any, use of leverage to increase their exposure to the market, some exchange-traded funds (ETFs) may use significant amounts of leverage, resulting in exposure that magnifies an index return two or even threefold. In an upward trending market spectacular gains may be achieved. In a downward-trending market, however, invested capital may disappear “quicker than the ticker.”

Note that the appropriate use of derivative instruments can sometimes attract criticism from those who have a less than perfect understanding of hedging. For example, an institution that enters into an interest-rate swap to place a ceiling on the total funding cost of its floating-rate debt issuance may be criticized by commentators if interest rates subsequently fall. Such second-guessing often assumes that the swaps were entered into for speculative purposes and fails to note the benefit that the swaps would have provided had interest rates instead risen. If the debt had already been issued, the institution would not suffer a loss, but rather would merely forego what would otherwise have been a prospective gain, since the loss on the swap would be offset by the reduced amount of interest paid. If the swaps were entered into in advance of, but with the expectation of, issuing debt,

it may still be considered to have been a prudent transaction if there had been reasonable grounds for expecting rates to rise. Of course, purchasing a swaption or an interest rate cap—whose description can be found, for instance, in Hull (2008)—might have been a better strategy, although these instruments require the upfront payment of a premium, which may have argued against their use.

Volatility Risk

Options are often priced using a quantitative model, one of the inputs to which is volatility. For example, a model used to price an IBM call option would typically use the volatility of the IBM stock price as an input. The higher the volatility, the greater the chance that the call option will be in the money at expiration and hence the greater the current value of the call. Since volatility cannot be determined in advance, the model may use *historical volatility*, which might not truly represent the current and future volatility, or it can use *implied volatility*, that volatility that would be required to have the model produce the current market price of other IBM options. The purchaser of the IBM call option is exposed to the risk that the level of implied volatility will fall, reducing the current value of his option and decreasing the chances that his option, if currently out-of-the-money, will expire in-the-money.³

Fixed-income investors, too, can be subject to volatility risk. The price of a bond with an embedded option such as a callable bond will vary as the value of the option changes, for example due to a change in the volatility of interest rates. Since the holder of a callable bond is effectively short the option, an increase in interest rate volatility, and a corresponding increase in the value of the call, will reduce the price of the callable bond. The bondholder is, therefore, exposed to volatility risk.

Large price swings in the market may also reduce liquidity and significantly impact trading costs.

Manager Risk

An investor in a mutual fund or other managed investment pool is exposed to the market risk of the vehicle in which he has invested, for example, an equity growth fund or a corporate bond fund. The *fund prospectus*

³In practice, implied volatility encompasses more than the current expectation of volatility. It includes liquidity and other premiums.

contains a statement of the *investment objectives* of the fund and the strategy that the fund manager intends to follow to achieve those objectives. It will usually also disclose the *benchmark*, which is typically a broad market index, although it may instead be a custom index, the return of which the manager will aim to match or exceed. The manager's portfolio will often differ from that of the benchmark in an attempt to generate a performance *alpha*, or excess return over the benchmark's market return or *beta*.

Manager alpha arises from *selection return*, since it is attributable to the specific securities in which the manager chose to invest, as opposed to the *style return* associated with the combination of market sectors that comprise the benchmark. As a manager's portfolio composition starts to differ from the makeup of the benchmark, the risk profiles of the two portfolios will start to diverge, with the manager's style return now being that of the sectors in which he has actually invested. An equity manager may seek incremental return by a variety of strategies, including purchasing only a subset of the securities in the benchmark or over-weighting particular market sectors. A fixed-income manager may speculate on the likely direction of changes in interest rates and set the duration of his portfolio accordingly by overweighting particular tenors on the yield curve. In both cases, the manager has chosen to *take a view* of the market with the aim of beating his target. To determine the success or failure of that choice would require the investment of that most precious of commodities, time.

Not all investors have an investment objective that measures returns against a benchmark. Some, particularly those who invest in hedge funds, have an *absolute return* objective. For example, rather than requiring a return that at least matches the return of the S&P 500, an investor may have a return objective of 15% per annum. Consequently, the traditional concept of alpha and beta is not relevant for such an investor.

Funds that allow investors to redeem their shareholdings often maintain a liquidity cash reserve to facilitate redemptions, particularly if some of the fund holdings are illiquid. If the funds do not, then in the event of redemptions, the manager would be forced to sell some of the fund holdings, possibly impacting the market price for the remaining fund holdings. The cash reserve can lead to fund underperformance against the benchmark, or negative alpha, in an appreciating market, since the reserve is usually invested in lower yielding treasury bills. Conversely, in a depreciating market a cash reserve may contribute to positive alpha. If a fund does not maintain a liquidity reserve, however, then redemptions during times of reduced liquidity can force the fund to sell assets at distressed prices.

In order to reduce the size of the cash reserve used to protect against *redemption risk*, hedge fund managers often impose a *gate provision*,

limiting the size and/or timing of redemptions. This can lead to increased risk for shareholders unable to exit the fund in volatile markets and is an example of liquidity risk for investors.

In some funds, for example, hedge funds, an investor may only have to pay a management fee if the fund shows a positive return for a given period. Investors who make use of more than one manager or who invest in more than one fund, however, can be faced with a situation where one fund has a positive return and so merits a management fee, while the other fund has a negative return of greater magnitude. While the performance of the latter fund does not merit a management fee, the investor may end up in the unhappy position of having a net negative return on his portfolio as a whole, while having to pay a management fee on some portion of his assets.

Benchmark Risk

Actively managed investment vehicles, for example, mutual funds, are often managed against a benchmark index. In other words, the stated investment objective of the fund is to match or exceed the return of the *bogey*, while subjected to a similar degree of risk. In practice, managers rarely beat the index. In many cases, however, this failure is because the return of the bogey is misleading. Broad market indices are periodically reconstituted, with some securities being added to the index in place of others. When a security has been added to an index, its market price will often increase as managers who are attempting to replicate the performance of the index will purchase the security. Similarly, when a security is dropped from the index, its price will tend to drop as portfolio managers eliminate it from their positions. In both cases, many managers will be trading at *ex-post* prices, while index performance will be calculated based on *ex-ante* prices. Managers with small portfolios relative to the size of the capitalization of the benchmark may find it impractical to replicate the entire index; trading costs associated with purchasing small lots of a large number of securities may be prohibitive. In addition, if the index contains issues with relatively small market capitalization, or small free float, it simply may not be possible for some managers to purchase the security.

Investors evaluating performance against a benchmark need to thoroughly understand the index methodology and the trading environment, including whether the benchmark returns are net of trading costs, whether the index level is recast in the event of company credit events or defaults, what the liquidity is for index constituents, whether the investment universe is truly investable, and the appropriate tax treatment.

Inflation Risk

Inflation, which we discussed in Chapter 1, is of concern to investors because it reduces purchasing power. When lending money, for example, by purchasing a bond, an investor needs to be compensated for expected inflation.

To mitigate inflation risk, inflation-indexed bonds such as U.S. Treasury Inflation-Protected Securities (TIPS) have been introduced. The principal on these securities accrues with inflation, and the coupon is a real, rather than nominal, return on the inflation-accrued principal.⁴

Correlation Risk

Modern portfolio theory, pioneered by Harry Markowitz, demonstrates the benefits afforded to investors by diversification. By investing in assets that are less than perfectly correlated, investors can optimize their expected return while minimizing their risk. Consider, as an example, two manufacturing companies, A—a producer of sunscreen products—and B—a producer of umbrellas. Suppose we were to invest in both A and B, each with an expected return of 10%:

Investment	Weather	Probability	Return (%)	Expected Return (%)
A	Sunny	0.5	20	10
	Rainy	0.5	0	
B	Sunny	0.5	0	10
	Rainy	0.5	20	

Assuming that the weather can only be sunny or rainy, with equal probabilities, then the expected return of each investment is 10%. Notwithstanding the positive expected return, however, if an investor were to invest in only A or B, he would have a positive return only 50% of the time. However, if he were to allocate his investment equally to A and B he would have the same realized return all the time, irrespective of the weather. By diversifying into assets that were not fully correlated—or, in this case, were inversely correlated with each other—the investor reduced the unsystematic risk in his portfolio. (A comprehensive discussion of modern portfolio theory and

⁴In practice, these bonds incorporate a liquidity premium in their price due to relative scarcity. For the nine months through June 2009, for example, just \$44 billion of the \$6.66 trillion of U.S. government bonds issued were TIPS.

efficient frontier analysis may be found in, among others, Elton, Gruber, Brown, and Goetzmann, 2006.)

Implicit in this kind of diversification strategy is the assumption that correlations are stable. Correlation risk can arise from the fact that the relationships between assets in a portfolio can change over time, especially during periods of market stress. For example, from 1991–1994 the correlation between international stocks and the S&P 500 was 0.3 to 0.4, suggesting that a combination of the two asset classes would provide diversification, while commodities had a slight negative correlation with the S&P 500. However, in 2008, the S&P 500 dropped 37% while the Morgan Stanley Capital International (MSCI) Europe, Asia, and Australia index dropped 45%, and commodities fell 37%.

Investors should understand, first, that correlation does not necessarily imply causation, that is, that a statistical relationship between two assets may simply be the result of chance, and second, that correlations—and relationships—may alter over time due to changes in investor behavior and macroeconomic and political trends.⁵

It is also worth noting that seemingly diversified investments such as *collateralized debt obligations* can be exposed to significant correlation risk. These instruments are bonds backed by cashflows from a pool of debt obligations, which typically are bonds, loans or other structured products.⁶ There may be a high degree of correlation among the underlying obligations. For example, if they consist of mortgage-backed securities, they may be concentrated in a geographic region and exposed to regional economic shocks. If the underlying cashflows are generated by corporate bonds, they may be issued by companies that operate in similar industries or trade in the same markets. It may be the case that a number of the obligations are guaranteed by the same insurance company; a failure by the insurance company to perform would likely not be confined to a single policy.

Synthetic CDOs can expose an investor to significant correlation risk. In a synthetic CDO, credit exposure derives not from the ownership of a debt obligation but from having written credit protection such as a *credit default swap* and can be much riskier. With a typical bullet bond portfolio, for example, consisting of 100 equal position sizes, a default by any issuer would result in a loss before recovery of 1%. Were 10 issuers to default, the portfolio would lose 10% of its value. However, with a synthetic CDO,

⁵While there may be a strong correlation between the number of bars and churches in any given town, for example, the number of each type of place of contemplation is more likely to be linked to population size than to the number of the other.

⁶A CDO that has other CDO cashflows as collateral is referred to as a CDO². The component CDOs may contain obligations of the same issuers.

that had provided credit exposure on the same 100 issuers, perhaps five companies could default on their debt obligations without the investor in the CDO suffering a loss. If a sixth company defaulted, the CDO investor could lose 25% of his principal. If 10 companies defaulted, the CDO investor could lose 100% of his principal. Because of the significance of even a small number of defaults on a CDO investor's position, it is vital that the portfolio be minimally correlated—and that correlations remain low. The reader is reminded of the principle of *adverse selection*, the possibility that the seller of a pool of assets may cherry pick his inventory to transfer underperforming assets and retain premium assets for himself.⁷ For a more detailed discussion of these topics, refer to Tavakoli (2003).

Political Risk and Sovereign Risk

The terms *political risk* and *sovereign risk* are sometimes used synonymously, but in fact they have different connotations. Political risk, also known as *country risk*, relates to the danger that an investor will suffer an economic loss due to non-market factors such as unfavorable political decisions, political instability, civil unrest, or macroeconomic or social changes. Examples of political risk include governmental nationalization or expropriation⁸ of property, such as that following the 1959 Cuban revolution, and Venezuela's 2007 nationalization of much of the oil industry and its 2008 takeover of the cement and steel industries.

Sovereign risk is the risk of economic loss due to a change in foreign reserve regulations with negative consequences for the convertibility of the foreign currency, or due to a sovereign government repudiating its debt obligations. Recent examples of sovereign risk include the Russian debt crisis of 1998, when the Russian government announced a moratorium on the repayment of some of its debt obligations, and the Argentine economic crisis when, in 2002, its Central Bank suspended the convertibility of its currency.

While often assumed to be associated with foreign investment, especially in developing economies, in practice both risks also apply to an investor's home market. In 1932, for example, the United Kingdom converted its 5% War Loan Bonds into new 3.5% bonds in what was in effect a sovereign

⁷An investor in a synthetic CDO might reasonably ask himself why the portfolio consists of the particular credit default swaps that it does. He may also wonder what prompted the writing of the protection in the first place, whether the price of protection on the debt is correctly priced in the market, and if there is model risk.

⁸Expropriation refers to the nationalization of an asset for which the compensation is deficient or nonexistent.




FIGURE 2.2 U.S. Treasury five-year credit default swap premium (cost of protection). Euro denominated.

Used with permission of Bloomberg L.P. and its subsidiaries. Copyright© 2010. All rights reserved.

default. The United States, meanwhile, with the passage of the Gold Reserve Act of 1934 in effect defaulted on its obligations by abandoning gold convertibility and changing the value of gold from \$20.67 per ounce to \$35.

More recently, at the peak of the 2007–2010 crisis, investors' confidence in the creditworthiness of even the biggest developed governments faltered, and the cost of sovereign credit protection soared. Figure 2.2 shows the premium on the five-year U.S. Treasury credit default swap (CDS), that is, the fraction of the notional—in basis points—the purchaser of credit protection would have to pay annually. On February 24, 2009, the premium reached a dizzyingly high level of 100 basis points, to decline subsequently to a value more than three times the pre-crisis level. CDS premiums on sovereign debt of other major industrialized countries exhibited a similar spike. To put this in perspective, before the crisis, the premium on a North America (mostly United States) investment-grade corporate CDS index hovered around 35 basis points, while the premium on a high yield index dipped to as low as 209 basis points in February 2007.

Note that credit default swaps on U.S. Treasuries are denominated in euros: A settlement in U.S. dollars, following a U.S. government default, would be of little use. Credit default swaps on other developed economies'

sovereign debt are usually traded in U.S. dollars and sometimes, for non-euro-zone countries, in euros.

Political risk, too, is not confined to less developed countries. In 2009, for example, senior secured debtholders in the United States were strong-armed by the government into accepting legally questionable restructurings of auto companies. Similarly, as discussed in Chapter 1, deficit financing has major implications for inflation, interest rate, and default risks.

Closely related to political risk is *legislative risk*, which arises from the possibility that a state or national government may enact laws that adversely affect the conduct of a particular business, or business in general. For example, laws may be enacted that increase the cost of providing health care or pension benefits to employees, increasing costs and resulting in a reduction in business profits. Changes in environmental, employment, and tax law could have major implications for the profitability of entire industries.

Useful resources for readers interested in political and sovereign risk include www.mkeever.com and www.transparency.org.

Market Sentiment Risk

No discussion of financial risks would be complete without mentioning market sentiment. Excessive market optimism and excessive market pessimism both pose substantial risks to an investor, when greed or fear lead market participants to ignore rational valuation models and instead to blindly follow the herd. As introduced in Section 1.1, the intrinsic value of an asset, particularly a financial asset, is usually predicated on an implicit expectation of the receipt of cashflows associated with the ownership of that security, whereas the market price of the asset is the extrinsic valuation level at which a willing buyer and a willing seller will currently execute a trade. In normal markets, with rational buyers and sellers, the value and market prices of assets should be congruent. From time to time, however, market participants collectively behave irrationally. During the dotcom bubble of the late 1990s, for example, assets were bid up significantly beyond the values that could reasonably be justified. Conversely, during the credit crisis of 2007–2010, many assets were trading at significant discounts to their intrinsic values.

Accounting principles state that a financial asset should be marked to market, that is, recorded on a firm's books using the market price of the asset. For a fixed-income investor who intends to hold a bond until maturity, the dichotomy between price and value is evident; a bond that does not default could be purchased at par, over time be accounted for at a price significantly below par, and at maturity be redeemed for par. A buy-and-hold investor, therefore, can have a very different view of value from an investor with a short-term perspective.

Market sentiment risk may impact the market as a whole, as in the recent financial crisis, when prices of almost all market sectors were affected, or it may be confined to a particular market sector or even to individual securities. For example, equity investors may find that a sector such as biotech may fall in and out of favor over time, with the market pricing in wildly differing earnings multipliers. While the trend may or may not be your friend, portfolio managers should act cautiously when value and price diverge. While such a divergence may be readily identified, the timing of any correction is difficult if not impossible to predict. History has demonstrated that investors who attempt to profit from an inevitable correction can be bankrupted before the market reverts to normalcy. As John Maynard Keynes is reported to have observed, “The market can stay irrational longer than you can stay solvent.”

Commodity Risk

In addition to financial instruments, many portfolios contain exposure to commodities, either directly or indirectly. Investors in airlines, for example, have significant exposure to oil prices, as do investors in transport and chemical companies. In fact, practically every company has some exposure to commodities, either directly through the raw materials that they use or indirectly through energy and distribution costs and the broad economic environment. It is often the case that the person responsible for managing a firm’s exposure to commodity risk is not a finance professional, and yet an understanding of financial markets is critical to the effective management of these exposures.

Commodities are often organized into the following groups: precious metals, base metals, energy fuels, petrochemicals, plastics, financials, freight, livestock and agriculture, and pulp and packaging material. They generally exist in finite quantities, and so are particularly influenced by demand trends. As the demand for oil increases, for example, the price will tend to rise. However, as the price increases, the supply of oil may expand as it becomes economically feasible to utilize previously untapped reserves and resources.⁹ The motivation for developing substitutes and the introduction of new technologies is also positively correlated with price.

Most commodities are priced in U.S. dollars and so will often see increased investment by overseas investors as the dollar weakens, thereby

⁹Oil reserves are those estimated quantities of oil that are highly likely (*proven reserves*) or likely (*unproven reserves*) to be recovered. *Oil resources* generally relate to petroleum quantities that may with time become recoverable, due to changes in political, technological, or other factors.

functioning as a useful hedge for U.S. dollar exposures. Commodities also prove attractive to investors seeking returns higher than those offered by fixed-income instruments when yields are low. As discussed in Chapter 1, during periods of rising inflation, commodities such as gold are considered to offer a good hedge against a weakening U.S. dollar and the deleterious effects of inflation.

Investors who wish to invest directly in commodities may do so through the use of exchange traded funds (ETFs), through the use of index futures, and through the use of a *managed futures fund* that is directed by a *commodity trading advisor* (CTA). The regulatory body overseeing commodities trading is the Commodity Futures Trading Commission.

2.3 OPERATIONAL RISK

The Basel Committee on Banking Supervision defines *operational risk* as “the risk of direct or indirect loss resulting from inadequate or failed internal processes, people and systems or from external events,” and so covers loss due to a breakdown in back-office systems and fraud. It also encompasses operator error, such as buy and sell orders incorrectly or mistakenly entered by traders. Since it relates in large part to the actions of individuals both inside and outside an enterprise, it is sometimes referred to as *people risk*.

An example of loss due to operational risk was the sell order for \$3.5 billion of shares in a company that was then valued at \$93 million, executed by a trader at Mizuho Securities in 2005. This order, the result of mistyping, should have been flagged by the order processing system but was not, resulting in a loss of \$335 million for the firm. In 1998, a Salomon Brothers trader leaning on his keyboard accidentally pressed the F12 key, triggering an instant sale of \$887 million of government bonds. Reportedly the traders at Salomon were not aware of the functionality of the F12 key, which, when pressed, replicated a previous transaction. Firms should have robust systems in place to guard against operational errors—and fraud—and, for example, to ensure as much as possible that unintended trades do not get executed.

Legal risk is the risk of loss resulting from exposure to legal issues, such as fines and penalties, lawsuits, and damages. Loss can result from a suit initiated by a counterparty seeking to avoid its obligations, or an insurance provider seeking to avoid payment on a policy due to a technicality. A firm should have legal policies in place and legal counsel employed to mitigate this risk.

2.4 WHAT RISKS ARE IN YOUR PORTFOLIO? HIDDEN HAZARDS

Many of the risks that we have just discussed will be readily evident when a portfolio is analyzed. Others may be less evident but may reveal themselves upon closer scrutiny. For example, while a bond portfolio is highly likely to be exposed to interest rate risk, it may contain mortgages and so be subject to call risk. An ABS deal backed by auto loan collateral may be assumed to have credit risk only to a nationally diversified pool of borrowers. However, if the collateral had been originated by a finance subsidiary of a large auto manufacturer—a *captive finance company*—the collateral for the loans is likely to be predominantly or exclusively vehicles manufactured by the auto company. If serious design flaws in those vehicles were to emerge, then the resale value for the collateral would be likely to plummet, with negative consequences for the subsequent performance of the bond. Meanwhile, an investor in a grocery chain may discover that a significant part of its business relates to its financing operations. An investor in a manufacturer of luxury goods may find that the majority of its profits—and losses—depend on how well it manages its foreign exchange flows. Other companies may experience dramatic swings in profitability from changes in commodity prices such as oil or from holdings in other businesses. In 2009, for example, Porsche reported earnings of €6.8 bn, whereas its revenues were only €3 bn, due to windfall profits related to its holdings in Volkswagen. The latter company experienced a short squeeze, as hedge funds that had previously sold its shares short scrambled to cover their positions, resulting in the share price rising fivefold to make Volkswagen for a time the world's largest company by market value.

When managing an investment portfolio, a portfolio manager must consider if there are indeed less obvious risks that need to be addressed. For example, when purchasing an interest rate hedge such as a swap from a bank, there is a risk that the bank will fail to perform on its obligations, resulting in a loss to the purchaser if the swap is in-the-money. While this exposure can be considered to be a form of credit risk, in practice it is referred to as *counterparty risk*. Note that counterparty risk can be *potential* if an option is out-of-the-money, due to the fact that the option could in the future become in-the-money.

The *cost to close* is how much it would cost to liquidate the outstanding contracts on a position. If a hedge position is in-the-money, then the cost to close is positive for the investor, and he has credit exposure to the counterparty. If the counterparty were to default, and the investor sought to enter into new hedge positions, the price for doing so would be the cost to close.

Counterparty risk is a particular concern when purchasing hedge protection in the form of a customized *over-the-counter* (OTC) derivative rather than a standardized listed one, for example, a *listed option* or a futures contract. In the case of listed derivatives, while there is still some chance of loss if the derivative is in-the-money, there exists an exchange known as a clearing house with the liability being shared among all the members of the exchange, each of whom is *jointly and severally liable* for the obligations of each member of the exchange. *Exchange risk* relates to the risk of the exchange and its members defaulting on their obligations.

Regulatory risk arises from the actions of a supervising body or regulator, rather than those of a government or legislature. Examples of regulatory risk can include changes in margin requirements for brokerage accounts, or in the rules relating to institutional sales. Firms are exposed to *compliance risk*, that is, the risk of loss, censure, and penalties resulting from violations of laws and regulations, and a failure to supervise the actions of their employees. In a brokerage firm, for example, if employees are not suitably licensed, or if managers neglect their supervisory duties, the firm may be subjected to financial penalties.

Investments may be adversely affected by news stories about a company that appear in the mass media or in the financial press. For example, the French mineral water company Perrier suffered a loss of sales and a reduction in its stock price in 1990 when minute traces of benzene were found in some of its products. This *headline risk* can also arise from incorrect news stories, and from rumors. For example, in 2008, a technical glitch led to the Bloomberg news service incorrectly reporting the death of Steve Jobs, the CEO of Apple, Inc. and led to a 5.4% intra-day drop in the company's stock price.

Some kinds of investments are exposed to *human behavior risk*. For example, the return to an investor in an MBS is dependent on the actions of the underlying borrowers: whether they repay their loans and, if they do, how long it takes them to do so. Investors in mutual and hedge funds can be affected by the actions of other shareholders, particularly when liquidity is low. At times, too, the actions of senior management in a corporation may be more self-serving than in the interests of the shareholders of the firm, for example, by pursuing acquisitions for self-aggrandizement rather than dispassionate and objective reasons. The returns to investors in asset-backed securities backed by mutual fund fees will be adversely affected by higher than anticipated redemptions by the shareholders. This human behavior risk can, in extreme cases, manifest itself as market sentiment risk, for example if a large group of equity holders seek to buy or sell at the same time.

Investors are subject to the risk of discontinuities and disruption in market trading, due to a trading halt on an exchange or a communications

breakdown. They also face the risk of an economic loss that is the consequence of an unsatisfactory or failed trade, the result of extreme market volatility causing a *stop order* to be executed at a level significantly different from the *stop price*, or a *stop-limit order* not to be executed.

U.S.-based equity investors may choose to restrict their holdings to U.S. companies, whereas many of those companies may derive substantial proportions of their earnings from overseas operations, and so are exposed to foreign exchange and increased sovereign and political risks. Similarly, investors may discover that the companies in which they have invested, though registered in one place such as Hong Kong, may be conducting their businesses somewhere else, such as mainland China. Investors in mutual funds may be exposed not just to the risk of the individual companies or bonds in the fund, but also to risks such as counterparty risk, liquidity risk, and *valuation risk*, which is the risk associated with the process that the mutual fund uses to price the assets and calculate the daily net asset value (NAV) for the fund.

When making a loan, or purchasing an asset, the investor should consider whether his interests and those of the borrower or seller are aligned or conflicting. When selling, there is usually an *asymmetry of information*, meaning that the seller often has more complete knowledge of the asset and its associated risks than the purchaser. When lending, by providing a mortgage, for example, the interests of the borrower and the lender will be aligned if the borrower has equity in the property—that is, has not borrowed the complete purchase price—since there is an incentive for the borrower to ensure that he performs on the loan. However, if the loan subsequently becomes underwater (i.e., the principal balance exceeds the current market value) due to a market downturn, then the interests of lender and borrower will diverge, as there is less incentive for the borrower to continue making loan repayments.

It is important for investors to be aware of both direct and indirect exposures in their holdings. They ought to conduct not just a quantitative risk analysis, such as understanding the duration of a bond portfolio or the price to earnings (PE) ratios of equities, but also a qualitative risk analysis. They should seek to understand if there is *key man risk* in a technology company, for example. In a pharmaceutical company, is there a pipeline of new products, or is a significant part of company revenues associated with products for which the patents will soon expire? Are demographic trends likely to affect consumer demand for a company's products? Is there a threat from the development and introduction of new technologies, such as might be faced by a company that manufactures incandescent light bulbs? Is there a risk that changes in environmental law will have adverse consequences for the firm?

When analyzing risk, a manager should seek to identify the exposures to the different types of security, geographic regions, industry sectors, and commodities that exist in the portfolio, and how the portfolio would be affected by changes in the behavior of borrowers or consumers and by macroeconomic and political changes such as a change of administration, or a change in tax law, for example.

Sometimes portfolio managers and investors may sacrifice long-term gains on the altar of short-term profits. For example, some institutional funds participate in securities lending to other firms or individuals who engage in short-selling, and yet it makes little sense to support the actions of those seeking to depress the price of an asset you own.

2.5 HEDGING MARKET RISKS

In this chapter we have summarized the common risks encountered by managers of portfolios of financial instruments. While many of these risks are clear to see, some, such as correlation risk, human behavior, or hidden geographic and sector exposures, are easy to overlook. A manager's first task is to identify all the obvious and hidden risks inherent in his portfolio. Moreover, he should continually reexamine the portfolio's risks to stay current with the assets' evolution and external developments. Having done so, how can the manager *deal* with these risks?

Whereas some risks, such as operational or political, can only be addressed by non-capital market means, such as ensuring smooth business functioning with adequate controls and, for political risk, judicious country investment decisions and diligent study of the political environment, market risks by contrast can often be substantially reduced through hedging.

Hedging means taking a market position that to some degree offsets an exposure present in the portfolio. In the purest sense, a hedge is not intended to make money but to prevent or limit a loss on an investment.

As an example, consider a manager of a U.S. equity fund with a small-cap exposure. Suppose the manager expects small-cap stocks across the board to suffer a temporary decline during the next month. Instead of selling the usually illiquid, small-cap names he holds and buying them back one month later, at great transaction costs, the manager can enter the short side of a one-month futures contract on a small-cap index, for instance the Russell 2000. Futures contracts on major indices are actively traded on exchanges such as IntercontinentalExchange (ICE) and Chicago Mercantile Exchange (CME) and are much more liquid than the collection of underlying stocks. If the small caps decline, as the manager expects, then the gain on an appropriately sized short contract should closely match the loss on the small-cap portion of

the fund. Of course, if the small caps rally, then the hedge will likely negate the potential gain. In order to keep the potential upside the manager may choose a *nonlinear* hedge, for example, a *put option* on the Russell 2000 expiring in one month. Such an option would protect well on the downside but, unlike a linear hedge, will allow the manager to keep much of the upside should his expectation prove false. In either case the premium paid for the option would diminish the fund's return.

This example illustrates a number of issues relating to hedging. The immediate question is, What should the notional of the futures or option hedge be? To answer this question we need to *quantify* market risks. Specifically, we need to know the fund's sensitivity to changes in the Russell 2000. For the purpose of hedge sizing, it is most convenient to express this sensitivity as *equivalent variation investment* (EVI) in the index, namely an investment of such a sum of money that its variations with small moves in the index equal the asset's value changes attributable to these moves. In Chapter 3 we will discuss in greater detail this and other measures of sensitivity and show how to calculate them.

A full futures hedge would have the exposure magnitude, as measured by EVI, as the notional. The manager may, however, decide to hedge only part of the exposure, say 50% of the EVI, to mitigate the impact of possible index decline while at the same time retaining an opportunity to capture some of the upside.

A put option could be bought on a notional equal to EVI, thus placing a floor under the potential loss without substantially reducing short-term profit-or-loss (P/L) volatility. Alternatively, the manager could choose to *delta-hedge*, in which case the notional will be the exposure divided by the option's delta. For small index changes, the performance of the delta hedge will not be much different from a linear hedge with futures (except, of course, that one is poorer by the amount of the option premium), but for a large move in either direction the delta-hedging strategy will deliver superior performance thanks to the option's *gamma*: In the case of a delta hedge one can view an option's premium as the price of gamma. As with futures, the manager may deploy only a partial hedge protecting less than the full exposure.

The fund is unlikely to hold every component of Russell 2000, and in exactly the same proportion as the index, and therefore the changes in the fund will not match those in the EVI exactly. The difference is called *basis* and is the cause of *basis risk*. In practice, if the fund has a significant number of typical small-cap stocks (say 10 or more), then the hedge is likely to work acceptably well for large shocks to the small-cap universe, but not necessarily for small changes. Note that basis risk was introduced in this example by the manager's attempt to reduce the cost of hedging the portfolio. Here, as

in many cases, increased exposure to basis risk is the result of seeking to decrease the consequences of exposure to *liquidity risk*. In Chapters 3 and 5 we will discuss how to assess the goodness of an index as a proxy for a particular exposure and we will delve into the subtleties of basis risk and attendant trade-offs in Chapter 6.

As the reader may have surmised from the foregoing discussion, market risk is the mirror image of profit: Risk is the “downside half” of an exposure while profit is the “upside half,” and a hedge seeks to limit potential loss at the cost of potential profit (*opportunity cost*). One could say that a perfectly hedged portfolio is a dead portfolio since a complete elimination of risk also eliminates profit potential (in fact, the portfolio would be steadily bleeding the cost of hedging).

In addition to curtailing the upside, a hedge can introduce novel risks such as counterparty exposure or the always unwelcome negative gamma. A corollary is that some exposures may be better left unhedged or should be hedged only partially; before putting on a hedge a manager should balance his tolerance for downside risk and P/L volatility against the additional basis risk, novel risks, and hedging costs, both direct and as foregone upside. When a decision to hedge has been made, the manager needs to decide on the instrument to use, whether it be linear, nonlinear, or a combination thereof, as well as on the extent of the hedge, namely the full exposure or only a portion.

Once in place, a hedge needs to be managed. This can mean changing the hedge size, substituting one hedging instrument for another, rolling (e.g., entering a new futures contract when an old one expires) or removing the hedge altogether based on market moves, macroeconomic conditions, the evolution of the asset, or changed risk tolerance. Managing hedges involves a considerable degree of discretion, which can be used by a skilled manager to enhance overall return while keeping market exposures low.

We will address quantifying risks and calculating EVI and other sensitivity measures in Chapters 3 and 5 before proceeding with a detailed discussion of market risk hedging in Chapter 6.

CHAPTER 3

Asset Modeling

Consider whence each thing is come, and of what it consists, and into what it changes, and what kind of a thing it will be when it has changed, and that it will sustain no harm.

Marcus Aurelius Antoninus, *The Meditations*

3.1 ASSET VALUE

What is an asset worth?

The answer is usually expected to be a number X of monetary units (U.S. dollars, euros, pounds sterling, yen, etc.) considered to be equivalent in some sense to the asset. This equivalence may be in the sense of *price*. For instance, an asset holder may be able to sell the asset to a buyer offering X dollars (to be specific) for it—this is the *bid* price—or an aspiring owner may have identified a seller asking the same amount—in which case X is the *ask*, or *offer*, price. If the asset holder is willing to part with the asset for $\$X$, then the asset is worth at most $\$X$ to the holder. Likewise, the aspiring owner who pays $\$X$ to become an actual one values the asset at $\$X$ or more. Finding a willing counterparty, however, does not mean that a transaction will be completed: The asset holder may conclude that the bid price is too low or the would-be owner may deem the ask price too high. In refusing to transact, they would rely on their assessment of the asset's *economic value*, that is, the economic benefit conferred on the asset holder.¹ The prospective seller and buyer would conclude that the bid or ask price does not reflect the asset's true worth.

Price as a candidate for the measure of an asset's worth can also enter through a consummated trade: If an identical asset changed hands most

¹ We limit our discussion to economic considerations, disregarding, for instance, sentimental value the owner can derive from an asset.

recently for \$X, should not then this *trade price* be an accurate measure of the asset's worth? This proposition has three potential flaws. One is observational: While liquid exchange-listed stocks may trade hundreds of times per second on the exchange, each trade's price and volume recorded and available in real time, less liquid securities may trade very infrequently or not at all; the "last trade price" may be hours, weeks, or even months old, or simply absent. In addition, no record is generally available for assets traded predominantly or exclusively over the counter.

A second potential flaw is the deviation of the observed price from the asset's economic value, that is *mispricing*. While in the "long run," prices are expected to revert to the assets' fundamental value, short-term price fluctuations are driven by news, supply and demand, investor psychology, and the idiosyncrasies of algorithmic trading systems. In addition, as will become clearer later in this chapter, the "economic," or "fundamental," value of an asset is a fuzzy concept: It is a function of the valuation method, the details of this method's assumptions and implementation, and the quality of input data. In other words, while there may be good and bad asset models, there is no single "right" one producing a "correct" economic value to which prices should converge. In fact, a good model should produce not a single number but a confidence interval of statistically significant valuations.

The observed price may additionally diverge from the economic value as a result of a *forced sale* or *buy*. For example, a distressed financial company may have to sell illiquid asset-backed securities to the highest bidder; an asset management firm may have to sell large quantities of securities to meet investor redemption requests, in the process depressing the asset prices; a *short squeeze* can result in spectacular price spikes. A dramatic example of a short squeeze occurred on October 27–29, 2008, when the share price of Volkswagen AG jumped 336% in two days, on a 2.5-fold increase in trading volume, only to fall 44% on the third, as illustrated in Figure 3.1. Volkswagen price and volume spiked following Porsche SE's unexpected announcement that it had raised its stake in Volkswagen from 35% to 42.6% of ordinary shares and held options for another 31.5%. A few days before the announcement, almost 13% of Volkswagen's common stock had been on loan, mostly for short sales. As short sellers scrambled to cover their positions in order to limit losses and avoid posting additional margin, the demand-driven price jump was compounded by the scarcity of common shares available for sale: Most of Volkswagen's common stock is closely held, Porsche's options' writers had covered their exposures by buying the underlying, and much of Volkswagen's remaining free float was held by index-tracking funds. These funds' objective is to replicate the performance of an index rather than maximize the return on investment and,




FIGURE 3.1 Volkswagen AG short squeeze. The price (top; euros) and volume (bottom) spiked on October 27–29, 2008.

Used with permission of Bloomberg L.P. and its subsidiaries Copyright© 2010. All rights reserved.

consequently, their strategy constrained them from selling holdings to lock in a gain.

Forced sales and buys are also a feature of index funds, as well as some actively managed funds that strive to minimize tracking error relative to the benchmark: When the benchmark index is rebalanced, such funds have to adjust their holdings to approximate the index’s new composition.²

Because of mispricing, the observed market price is usually a poor indicator of an asset’s intrinsic worth. In fact much of the money management industry is predicated on identifying and exploiting asset mispricings.

² Note that many index maintainers, most notably Standard & Poor’s and Russell Investments, change their indices’ composition on an ad hoc basis, usually with a one to five days’ notice. S&P replaces members of most of its U.S. equity indices whenever warranted by market developments or corporate actions. Russell, by contrast, reconstitutes its equity indices on a preset date once a year. Between the reconstitutions, however, Russell *removes* stocks that cease to meet the index membership criteria and adjusts the remaining stocks’ weights. Consequently, on most days of the year, the Russell 3000 index, as an example, would have fewer than 3,000 members. See www.indices.standardandpoors.com and www.russell.com/Indexes for more details.

Finally, a third problem with using the last trade price as a measure of an asset's worth relates to the word "identical": There may not be another asset with exactly the same characteristics. This is especially true for many asset-backed securities, including mortgage-backed, as well as private placements and real estate. For such unique assets market participants often use sales of similar assets as a guide, but accounting for the differences may be subtle and imprecise.

An alternative to observing an asset's price is to equate its intrinsic value with the *present value* V of the asset's future cash flows: This is the economic benefit the owner of the asset derives from holding it. This valuation approach is a special case of the *discounted cash flow* (DCF) analysis and is the main subject of this chapter.

If there were a way to determine V "exactly," then all the trades would be executed at this "correct" price. But since the economic benefit acquired in a purchase would exactly equal the price paid less transaction costs, and the money received for selling an asset would equal the foregone benefit less transaction costs, trading activity would only result in frictional losses and would therefore cease. Luckily, the economic universe is unlikely to end by entropy death for three main reasons. One is that participants in the economy may wish to purchase assets in order to *add value* by combining them with other assets and with factors of production. For example, a jet engine manufacturer may buy titanium to combine it with other materials and the expertise of its labor force (a factor of production) to produce an engine whose economic value (determined, for instance, as the present value of lease payments over the engine's life time) will exceed the combined cost of ingredients and labor. Specifically, the contribution of the purchased titanium to the engine's value added will exceed the price gain that the manufacturer could obtain by reselling the metal by the time the engine is ready. The titanium producer, however, cannot capture this value added because he cannot manufacture jet engines any more than the jet engine manufacturer can mine ore and (efficiently) extract titanium.

Similarly, a mortgage originator may sell mortgages to a structured products desk, which will then pool them into asset-backed securities and sell to other investors. By combining the bought mortgages with its other assets (skilled employees, business relationships, specialized expertise), the desk adds value relative to the worth of the underlying loans.

A money manager may also acquire a financial asset not for its own benefit but to combine it with other assets in the manager's portfolio in order to improve its *risk-adjusted return*. This kind of asset would typically be negatively correlated (or, in some cases, be effectively uncorrelated or even slightly positively correlated) with the portfolio and serves as a *diversifier*. It would reduce the portfolio volatility and, even if the portfolio's

expected return remains unchanged, would therefore increase risk-adjusted performance measures such as the Sharpe ratio, the Sortino ratio, and the Treynor ratio, which we define in section 5.2. For example, the manager may add commodities or gold to an equity portfolio not for their own sake but as a hedge against inflation: High expected inflation will reduce the numerator in the Sharpe ratio through the increased risk-free rate, but this can be offset by the concomitant reduction in the volatility in the denominator (the correlation between the returns on equities and commodities is likely to become “more negative” in an inflationary environment). As different managers have different portfolios and different *utility functions* (Campbell and Viceira, 2002), thus they may wish to trade assets at their respective perceived values

A second reason why trading activity will not grind to a halt is more relevant to this book and, fundamentally, is more important for the existence of active capital markets: There is no “correct” way to determine a unique present value of an asset’s future cash flows. The obstacles are twofold: (1) For most assets future cash flows can only be estimated, often with great uncertainty, and (2) the choice of the interest rates used to discount these cash flows is a matter of judgment and the rates are therefore subject to uncertainty. (The exceptional assets for which future cash flows are known with practically complete certainty are sovereign securities in national currency issued by major developed governments.)

Finally, an asset holder may wish to sell in order to use the proceeds for consumption or to redeploy them in a potentially more profitable way.

In this chapter we outline the principles of valuing an asset using DCF and demonstrate how to calculate an asset’s sensitivities to market factors given a valuation model.

3.2 FINANCIAL MODELS

Quantitative financial models have received much vitriol since the onset of the financial crisis of 2007–2010 and have even been blamed for it. This criticism has been misdirected. While the causes of the crisis are multiple and complex, with roots in some cases going back decades, one contributing factor was the *fashion* in which models have been mindlessly applied and blindly trusted. As with most intellectual developments, the problem lies with the users, not the modeling as such.

Humanity invented mathematics because it makes the formulation of ideas and their logical development easier—the authors invite readers with scientific background to state the Maxwell equations and prove, without recourse to differential and integral calculus, that they permit a solution

of propagating electromagnetic waves (light). The language of mathematics provides an efficient means of formalizing observations, knowledge, and intuition and, combined with the machinery of logic, of developing these inputs and deriving their logical consequences. The utility of this formalism, however, does not relieve its user of the responsibility for the soundness of the inputs and logical steps and for ensuring that they are up to date with the evolving reality. The most important rule, however, is that *the user should know in advance the intrinsic limitations of any quantitative model and should not draw conclusions extending beyond these limitations*. Any mathematical model is by necessity limited: Out of the continuum of facts and relationships the modeler must select a discrete set, and no investment or hedging decision should be made solely on the basis of the outputs of any particular model.

Another important rule is to think freely and not to be straitjacketed by standard, generally accepted categories, classifications, and nomenclature. A good modeler should have a fresh, unprejudiced look at the assets he intends to model and the purpose of the modeling exercise (valuation, investing, risk management, performance attribution, reporting, any combination of these) and, building on the collective experience of the quantitative finance community, think of the best way to approach his task. Confining oneself *a priori* to a particular kind of model, even with a fancy name, can distract from the nature of the problem and make the modeler overlook a novel or unique approach particularly well suited for his specific undertaking.

Elements of a Financial Model

With all the caveats of the preceding paragraphs, a well-designed and robustly implemented model provides a useful conceptual framework and an important reference point for appraising an asset's value and risks. In Figure 3.2 we summarize the main logical elements of a quantitative valuation model. The starting point is the basic *concept*, which should be rooted in sound economic principles but may also incorporate behavioral, political, and other considerations. The first task is to understand the security, its *inherent* properties and characteristics. Next, the modeler needs to identify a set of *most important* economic, market, credit, liquidity, behavioral, political, and other *external* factors that affects the value of the security as well as to devise these factors' *quantifiable measures*. Note the emphasis on "most important": in principle all conceivable factors affect the value of any security. The art of creating a good model involves identifying the ones that affect the security's value most directly while avoiding redundancy and not missing an important influence. Deciding on quantifiable measures is usually straightforward for the "quantitative" (e.g., economic and market) factors:

concepts involves a considerable element of art, designing a formalism is a solid engineering job.

With the formalism articulated, the model needs to be *implemented*, usually as a computer code in one or more programming languages. The implementation comprises not only the coding of mathematical formulae but also the creation of the necessary data feeds and organizing storage for the model's inputs and outputs. The choice of the specific programming languages, computational packages, databases, operating systems, and hardware platforms is usually influenced by the company's existing infrastructure, technology support, and the model developer's experience and preferences. In making this technological decision, one should consider the robustness and performance characteristics of the programming infrastructure, the potential development speed and ease of maintenance of the product and, very importantly, the *scalability* of the resulting implementation. In the past 10 to 15 years, many financial companies were burned on outsourcing application development to small companies, often one- or two-person enterprises located in Eastern Europe. Ostensibly cheap, such products soon revealed hidden costs when it became clear that the small firms did not have the resources to maintain and enhance their product in a timely manner; worse, some enterprises simply went out of business, leaving their customers with outdated products they could not maintain or upgrade.

Many financial models make use of a number of component "sub-models," often developed by different groups or individuals in their respective areas of expertise. A mortgage model, for example, may build on an interest rate model, a prepayment model, and so on, each of which may also be used by models developed to value other kinds of security. This approach helps ensure a consistent treatment of common market factors across security types and reduces model development effort through eliminating duplication. It does, however, add an extra burden of ensuring that the sub-models can communicate with each other and the final model. In larger organizations with extensive model development, it is advisable to establish a set of overarching specifications and protocols to minimize sub-model integration efforts and reduce the chances of miscommunication.

Model Evaluation and Maintenance

Needless to say, the implemented model should be thoroughly tested. The specifics of the tests would depend on the nature of the asset and the peculiarities of the model, but two broad types of test can, and should, be applied to most financial models. One is the *extreme value test* where some or all of the inputs are set to atypically, and usually unrealistically, high or low values. In such limiting cases, the asset value can often be calculated analytically,

either exactly or to a good approximation, and the model should reproduce this special result. For example, setting the default probability to nought in a corporate bond valuation model should result in a value equal to the sum of discounted coupon and principal payments, which, given a yield curve, can be easily calculated exactly. Extreme value tests do not necessarily probe the accuracy and predictive power of a model but rather its internal consistency and the correctness of implementation.

The other test type is *historical*. It consists in feeding the model actual historical data, especially from periods of market turbulence, and comparing the model output with observations. This type of test does appraise the quality of the model's underlying assumptions and its ability to reflect the market environment and respond to its changes. It is very important to conduct historical tests *out of sample*. This means that the testing period should come after the period used to calculate the model's parameters such as volatilities or correlations. In-sample tests are liable to exhibit spurious predictive powers since they test precisely the circumstances used in the model's calibration.

What are the attributes of a good model? Naturally, it should produce values and sensitivities corroborated *ex post*: Even if the current asset price diverges from the model valuation, the price should evolve toward the valuation with time. This predictive power strength criterion, however, is of limited use for illiquid securities with prices that are rarely, if ever, observed. However, the *changes* in the asset value (and price, if available) should be explainable in terms of changes in the market, economic and other factors via the sensitivities produced by the model regardless of the asset's liquidity. To a degree this is a test of the model's *self-consistency*: The change in valuation should be consistent with the sensitivities calculated using the same methodology.

Some qualities should be built into a good model from the beginning, at the stages of conceptual design and implementation. As already mentioned, the model should have a solid economic foundation. In addition, it should be *frugal*, in the sense that it should produce accurate results with the minimum of inputs, the smallest possible number of assumptions, and without undue complexity of the conceptual underpinnings and formalism. A good model should grasp the *essential* aspects of an asset while eschewing excessive refinements that usually do little to increase the model's usefulness. From the risk management perspective especially, it is more important for a model to respond agilely to abrupt changes in the economic environment than to calculate the *theoretical* value to the eighth decimal point or incorporate the effect of an arcane factor that changes the value and sensitivities by at most a few percentage points. One of the undesirable side-effects of excessive complexity is that it makes models *unstable*, that is, prone to large valuation

swings in response to small changes to the inputs, and therefore unreliable; over-engineered models usually perform worse than their simpler robust cousins.

At the same time, the factor set should not be stripped down to the point of missing out on important external influences, in other words it should be *complete*. In section 4.7 we describe a test that would flag the lack of completeness.

A special danger lurks in using interdependent factors, a phenomenon known as *multicollinearity*. Models suffering from this problem are particularly insidious, as they can create a false sense of security by producing high statistical goodness-of-fit measures, such as the coefficient of determination, also known as the “ R^2 ” (see Appendix B for definition), while their outputs are unstable and imprecise. In addition models afflicted with multicollinearity do not allow one to distinguish between the individual impacts of different factors.

As the macroeconomic environment and market conditions change, so should the models. The creator and user of a model should continually monitor the adequacy of its assumptions, inputs, and implementation in light of the changing circumstances and make appropriate modifications.

Finally, a finished model should be reviewed and tested by an independent auditor. Most large and medium-sized financial services companies have internal model verification and risk architecture groups for this purpose; smaller companies can hire one of a number of consultancies that perform “sanity checks” on financial models. Even banks with internal model auditing resources might consider contracting an external reviewer for a fresh look or specialized expertise.

Input Data

The output of even the best model is only as good as the data it receives. Figure 3.2 illustrates the classes of input data used by most financial models. Practically all of them use some kind of market inputs: interest rates, stock prices, equity and fixed income indices, credit default swap premiums, implied volatilities, credit spreads, and so forth. Many models also rely on macroeconomic data such as money supply, unemployment rate, sentiment indicators, and others discussed in Chapter 1. Models also need to know the data specific to the particular asset valued. This could range from static descriptives, such as the strike price and expiration date for a plain vanilla option, to sophisticated characteristics of a complex asset-backed security. These characteristics can change with time and be difficult to obtain in a timely manner, an example being the average prepayment speed of a pool of loans supporting an MBS. Finally, some models, especially those used for

active management, may require subjective inputs that reflect the manager's views. This could be the manager's assessment of the probability of an interest rate increase at the next Federal Open Market Committee meeting, the expectation of the passage of a piece of legislation, the anticipation of the resolution of a legal suit, or the strength of a sell or buy conviction expressed as an integral number between -5 and $+5$.

Even data from established and generally reliable sources may contain errors or may not be delivered at the scheduled time. A good model should incorporate data-checking modules that identify missing, clearly erroneous (e.g., invalid dates, negative prices, negative volatility) and suspicious (e.g., outliers) inputs, should inform the appropriate parties, usually via email, and should follow a contingency protocol prescribing the use of either stale data, data proxies, or approximations in lieu of the missing and dubious data points. The output in such cases should be clearly identified as being the result of a model run with deficient or stale data.

Some kinds of input, for example, the constituents of a mutual fund, are available only with a time delay of months or even quarters, while others may be obtainable in a timely fashion but at a substantial cost not justified by the benefits derived from these data's expeditious receipt. An example of the latter is real-time stock prices, provided by exchanges at a substantial cost, whereas the prices can be obtained at no cost with a time delay of typically 20 minutes. It would probably make little sense to pay for live prices to feed a model that is run only once or twice a day: The manager is likely concerned with daily and longer-term changes, not sub-hour fluctuations. However, the expense of purchasing real-time feeds would be fully justified, and in fact unavoidable, for an intraday trading system.

For a model to be useful, its designer should take the availability and timeliness of inputs into consideration and not rely on unrealistic expectations of extensive and punctual data delivery. These issues should be addressed at the stage of devising the model's formalism and be reflected in its implementation. In other words the constraints of data acquisition can, usually should and often do influence a model's details, as indicated with a black broken arrow in Figure 3.2.

Example: MBS Pass-Through

As an example let us consider an MBS pass-through security. These securities are pools of mortgage loans with similar characteristics, similar in the sense that loans backing a given pass-through would typically have about the same coupons, expected time to maturity, credit quality, outstanding balances, and so on. A given security would be characterized by such measures as weighted average coupon (WAC), weighted average time to maturity

(WAM), duration (sensitivity to interest rate moves), and convexity (duration's sensitivity to interest rate moves). Other important properties would be the issuer (one of the government-sponsored enterprises—GSE—or a private mortgage conduit), geographic concentration, and the size of outstanding loan balances. An especially important characteristic of all mortgage-backed securities is the *prepayment speed*, namely the rate at which borrowers retire the principal in excess of the scheduled amount.⁴

The ability to prepay the principal is a valuable call option owned by the borrower. From the investor's point of view, by contrast, prepayments can drastically, and most likely unfavorably, alter the pattern of cash flows and consequently significantly affect a pass-through's value. One of the most important single determinants of the prepayment speed on fixed rate mortgages is the prevailing level of interest rates: If the rates decline, then borrowers tend to refinance more than normal to reduce monthly payments and MBS investors are forced to reinvest the repaid principal at lower rates, a hazard known as *contraction risk* because additional refinancing shortens the expected life time of the mortgage pool. When interest rates rise, the borrowers, by contrast, refinance less than usual and investors receive slower than projected cash flows thus missing out on the opportunity to reinvest at higher rates. This phenomenon gives rise to *extension risk*, so called because the pool's expected lifetime is now longer. But economic factors other than interest rates affect the refinancing rate, and therefore the prepayment speed, as well. For instance, a shrinking economy, high unemployment, tight credit, and, especially, collapsing home prices can stymie refinancing even when rates are low—this is what happened in May–June 2009 when historically low mortgage rates failed to spur refinancing activity.

Two other key characteristics of mortgage-backed securities are the *delinquency rate* and the *foreclosure rate* of underlying loans. Normally, delinquency and foreclosures are less consequential than prepayments but in times of crisis they can become crucial factors in MBS valuation. This certainly was the case during the crisis of 2007–2010: according to the Mortgage Bankers Association of America (MBAA) (www.mortgagebankers.org), at the end of the first quarter of 2009, the seasonally adjusted delinquency rate (i.e., the percentage of loans at least one payment past due but not in the process of foreclosure) was 9.12%, the highest since records started in 1972 and 2.77% higher than one year earlier, when it measured 6.35%. The percentage of loans in the foreclosure process at the end of the first quarter of 2009 was 3.85%, a 1.38% increase from the first quarter of 2008, both

⁴ Analysts often distinguish between *prepayment*, the paying off of the full outstanding principal balance, and *curtailment*, a periodic or irregular payment in excess of the schedule of an amount smaller than the outstanding balance. This distinction is immaterial for our discussion and we will use “prepayment” for both phenomena.

the percentage and the annual change being record high. A shocking 12% of Florida mortgages were in the process of foreclosure at the end of the second quarter of 2009, according to MBAA. The rate was substantially higher for adjustable rate mortgages (ARMs): 19.97% for prime and 43.84% for sub-prime loans. Along with high unemployment, a major contributing factor to this spike in foreclosures and delinquencies was a precipitous decline in home prices: The Seasonally Adjusted U.S. National Values S&P/Case-Shiller Home Price Index (www.homeprice.standardandpoors.com) fell 31% between its peak in the first quarter of 2006 and the first quarter of 2009, leaving many homeowners underwater, that is, with an outstanding mortgage balance greater than the home value. Even many of those “above water” did not qualify for a refinanced loan, since lenders had tightened their rules and, in addition to stricter credit prerequisites, were now often requiring a loan-to-value (LTV) ratio of 80% or less; this meant the value of the home had to be at least $0.20/0.80 = 25\%$ higher than the outstanding balance—a dramatic change from the pre-crisis practice when LTVs of 105% and even higher were not uncommon (i.e., the home-buyer could borrow money for closing expenses). This inability to refinance deprived homeowners, cash-strapped in the midst of a recession, of an opportunity to lower their monthly payments in a low-rate environment, thus leading to missed payments and, ultimately, for many, foreclosure.

The mortgage market, more than others, can be directly affected by political developments. For example, in the United States the sociopolitical agenda of extending home ownership to lower-income groups led to the creation of government-sponsored enterprises, such as Fannie Mae, Freddie Mac, Ginnie Mae, and Federal Home Loan Banks, which enjoy, or did until 2008, explicit or implicit government guarantees and have dramatically affected the U.S. home-loan industry. The U.S. government sometimes attempts, often for political reasons, to influence mortgage activity by modifying the charters of GSEs, as it did in November 2008 when it increased the loan limit for “conforming” mortgages (i.e., those that satisfy the GSEs’ underwriting criteria).⁵

Taking all of this into consideration, a modeler attempting to value a pass-through would identify WAC, WAM, duration, convexity, prepayment speed, delinquency rate, and foreclosure rate as some of the security’s most important characteristics.⁶ He will then determine that the prevailing

⁵ For a more detailed discussion of GSEs and the MBS market in general the reader may consult, for instance, Fabozzi (2005).

⁶ For private label pass-throughs (i.e., those issued by entities other than GSEs), the modeler will also consider the average credit score of the borrowers. We do not delve into this and some other important aspects of MBS valuation, such as the effect of geographic concentration, in this simplified illustrative example.

interest rates are the principal market influence and choose the specific rates measures, say the London Interbank Offer Rates (LIBOR) for tenors up to 12 months and three-month LIBOR swap rates for tenors greater than one year. Next, the modeler may select a broad equity index, such as the S&P 500 in the United States, to gauge the broad state of the economy and investor sentiment. Knowing the significance of the labor market condition for prepayment, delinquency, and foreclosure rates, he would choose an employment statistic, for instance, the unemployment rate compiled by the Bureau of Labor Statistics and available at the Bureau's website (www.bls.gov).

Before the modeler can apply all these insights to value the pass-through, he needs to decide on a number of additional points. First, he needs to specify exactly how the external factors affect the security and its characteristics. One way to do this is by *regressing* the pass-through's characteristics on historical values of macroeconomic factors. For instance, one could regress historical prepayment rates on historical interest rates; one could also regress historical delinquency and foreclosure rates on unemployment rates and a broad equity index. The modeler needs also to determine the type of regressions to use: equal-weighted or with a historical decay, standard least squares or robust, the frequency of the data used and the length of the historical window (we discuss the technical details of regressions in Appendix B). These decisions depend on the nature of the asset and inputs used in its valuation. Much of the data used in an MBS model are available only with monthly frequency, so that in our example the modeler would perform regressions of monthly data series over a period commensurate with the WAM of the pass-through (preferably somewhat longer), probably with a mild time decay with a half-life of five to 10 years. The standard least squares method would usually suffice, but if the regression period involves unusual volatility, then a robust regression might be desirable. Ideally, an analyst should experiment with a range of regression types and decide on the ones with the greatest historically predictive power. Since different methodologies have their advantages and disadvantages, one might settle for running a number of regressions concurrently and comparing the results. Of course, at every stage the modeler should apply common sense and make sure that the relationships derived from historical regressions are sensible. For instance, a sharp increase in the unemployment rate, especially in the environment of high interest rates, should depress the value of an MBS pass-through. A model predicting otherwise is almost certainly flawed.

Second, the model creator needs to select a method of *projecting* the established relationships into the future. This can be done either through a "single-path" scenario or stochastically by using Monte Carlo simulations, as described in section 3.7. Finally, the modeler needs to find a way to incorporate *qualitative factors*, such as potential and actual legislative actions

and market and economic sentiment, on the pass-through's characteristics. For instance, the increase of the loan limit on conforming loans in November 2008, *ceteris paribus*, would have increased the prepayment rate in the pass-through's underlying pool of mortgages *if* the pool contained a significant number of loans with outstanding balance between the old and the new limits. In this case, and in general, the treatment of qualitative factors involves a considerable degree of discretion and subjectivity.

At this point, the modeler is ready to proceed with *implementing* the MBS pass-through valuation model. The programmatic details of the implementation, such as the programming language, the database, and the operating system would be determined, as discussed above, by the company's existing infrastructure and the developer's previous experience. The developer may use standard packages for statistical calculations, such as Matlab, SAS, and IMSL library. The implemented model should, of course, be subjected to thorough tests, including the historical and extreme value tests described earlier in this section, as well as independent audit.

Of course, this example is grossly oversimplified and serves only as an illustration. An actual model would take many more factors into consideration and be much more sophisticated. For instance, the prepayment rates are usually calculated by a separate *prepayment model* that takes into account some important factors we omitted for simplicity: asset seasoning (time since origination), seasonality (the variations in prepayment activity throughout the year), borrower burnout (a slow-down in prepayment activity, even when rates are low, characteristic especially of seasoned assets), and so on. The example, nonetheless, illustrates the basic elements present in a realistic quantitative financial model.

Model Risk

The foregoing general discussion highlights the importance of *model risk*. Creating any model involves discretionary decisions at every step, from the original concept to the implementation, and distinct equally "good" models can produce materially differing valuations. Naturally, the room for discretion and, consequently, uncertainty increases with the complexity of the asset.

In addition, the implementation stage is prone to simple coding and data flow errors. Often these lead to nonsensical results and are therefore easily detectable. The more insidious mistakes, however, can affect the model output in significant but not obvious ways. Errors of this type are difficult to detect, and the developer should exercise extreme care to avoid them. To a degree he can be assisted in this task by the use of a rigorous

object-oriented programming language with strong data type checking, such as C++: The language itself will disallow a number of common programming errors, and its object-oriented features will facilitate conceptualizing the code flow. Data flow errors can be kept in check by using a relational database with a well-designed normalized data structure. This kind of data storage will disallow, for instance, the insertion of an entry for a security with a misspelled identifier, a negative price, or an invalid date. Finally, extreme value tests can flush out some implementational errors.

It is convenient to group the sources of model uncertainty into four loose categories:

1. Conceptual
2. Formalism approximation
3. Implementational
4. Numerical

We touched upon the implementational errors in the preceding paragraph. These are true mistakes that can be avoided in principle but would inevitably crop up in any realistic project.

The *conceptual* uncertainty stems from the existence of distinct but equally valid ways of looking at a security and its value drivers. For instance, in the above MBS pass-through example, one model might put greater emphasis on the general economic conditions as a driver of the prepayment speed, while another may focus on employment statistics. In general, neither would necessarily be “better,” but they will produce differing valuations. (A significant difference between the two values, assuming the absence of implementational errors, may indicate that one of the two factors was more important during the time period used in calibrating and testing the model. In this case, a deeper look might be warranted, and one of the two approaches might turn out to be more appropriate for the period of interest and the task at hand.)

The *formalism* uncertainty reflects the multitude of acceptable ways to express the same concept in concrete mathematical and statistical terms. The ambiguity can arise, for instance, from the choice of the details of historical regressions (data frequency, historical window, regression type), specific model factors (S&P 500 or Russell 1000, Barclays Capital—formerly Lehman Brothers—or Merrill Lynch fixed income indices, NIKKEI 225 or TOPIX), the source of foreign exchange rates or LIBOR swap rates, the choice of credit default swap price provider, and so on. Since various formalizations are expressions of the same underlying conceptual principles, the corresponding differences in valuation should be insignificant. In fact, a

material effect on the output value of a formalization choice is probably a sign of a problem.

The *numerical* uncertainty derives from the purely computational limitations of numerical algorithms and is unrelated to the soundness of a model's underlying assumptions or the robustness of its formalism. In modern quantitative finance models, the numerical error is usually dominated by conceptual and formalism uncertainties. The model designer and user should, nonetheless, be aware of the ever present numerical fuzziness and be alert to its potential of becoming important. Below we list the main sources of numerical error and refer the reader to Press, Teukolsky, Vetterling, and Flannery (2002) for a detailed discussion.

The *roundoff error*, caused by the finite precision with which computer hardware can store numbers ("machine accuracy"), is rarely a problem in financial valuation models implemented on modern machines. One way to ensure it does not become an issue, even in especially intensive calculations, is to avoid using "single-precision" floating point numbers. In C and C++ programming languages, this means always using double (or even long double) and never float data types; it is not easy to conceive a financial problem where the additional storage requirement will exhaust an average modern computer's memory. The *truncation error* is similar to the roundoff error but arises not from the hardware limitations but rather from the approximations inherent in the numerical evaluation of mathematical functions, integrals, and the like. Just as the roundoff error, the truncation error is unlikely to be a cause of concern when valuing assets and can usually be warded off by requiring increased, even excessive, accuracy in such calculations: The high speed of modern computers will readily allow this. Roundoff and, especially, truncation errors can become important when calculating *sensitivities*, an issue we discuss in section 4.6.

A kind of numerical error that is more likely to afflict a financial model relates to calculating small differences of approximate numbers of similar magnitude, a situation arising in calculating asset sensitivities. To illustrate this phenomenon of *error amplification*, consider an exceptionally (and unrealistically) accurate model that can calculate an asset's value to within one part in 100. That is, if this model values an asset at \$1,000,000, then the combined effect of all the conceptual, formalism, and numerical uncertainties is \$10,000 or less (we assume the model to be free from implementational errors). Now suppose an analyst wants to study the sensitivity of this asset's value to changes in the S&P 500. To ensure he probes only the linear response, the analyst shifts the S&P 500 level by a small amount, say 1%, and reruns the model. If the elasticity (defined in equation (4.14)) of the asset with respect to the S&P 500 happens to be unity (i.e., the percent change of the asset value equals that of the index), then the model will produce

a valuation change of \$10,000, a number completely swamped by model uncertainties. Even if the asset is leveraged such that the elasticity is two and the valuation change is \$20,000, then the model error level will be about 50% of the change, much larger than the 1% uncertainty in the valuation. Of course, one lesson is to use bigger factor shifts to make sure that the valuation change exceeds the model error by a large multiplier. A shift too large, however, would be of little practical use if one is interested in the effect of small, especially intraday, market moves on the asset's value. The problem gets worse if one converts the valuation difference into a delta by dividing the valuation difference by the size of the small factor shift. We will revisit this issue in section 4.6; suffice it to mention here that there exists an optimal shift size that minimizes the sensitivity's error.

Another numerical error that a model creator is likely to encounter in practice stems from the finite number of trials N used in a Monte Carlo simulation. The fractional numerical error of a Monte Carlo simulation decays approximately as

$$\text{Monte Carlo simulation numerical error} \approx \frac{A}{\sqrt{N}}. \quad (3.1)$$

This is a one standard deviation error estimate corresponding to a 16% confidence level, that is, a confidence interval at a 5% or 1% confidence level would be a multiple of this. The prefactor A depends on the details of the model and the calculated quantity; it is typically somewhat greater than unity but rarely exceeds 10, so in order to be conservative one might set $A = 10$. Equation (3.1) implies that one needs at least 10,000, and possibly as many as 1,000,000, trials to keep the simulation error under 1%. This number of trials can be prohibitively expensive computationally when valuing complex instruments that require extensive calculations at each Monte Carlo pass, and, consequently, the achievable numerical error of a Monte Carlo-based model is often comparable to or exceeds conceptual and formalism uncertainties. The error in equation (3.1) is amplified in calculating sensitivities, as described earlier in this section, and the model user needs to ensure that N is large enough to produce meaningful results. In addition to evaluating equation (3.1), one should also vary N and compare the resulting outputs. When plausible, one should confirm that the output change resulting from the doubling of N is smaller than both the required accuracy and the size of other model uncertainties.

One more source of uncertainty peculiar to Monte Carlo arises from the discretion in the choice of the underlying random number generator. Not only can the choice of the generator *type* affect the results, but so can the selection of one of the many *streams of random numbers* a given generator can produce. It is always desirable to test a model with a number

of distinct generators and streams; the differences in the outputs should be economically immaterial. This generator-specific bias partially cancels out in calculating differences, and one should therefore use the same stream of random numbers in computing asset values for shifted inputs as for the base valuation. In practice this means that a model user should use the same “seed,” the number or collection of numbers used to select a particular stream produced by a random number generator.

We mention without going into details that there exist *variance reduction techniques* that allow one to achieve higher numerical accuracy than standard Monte Carlo for the same number of trials—that is, these techniques make the numerical error decay faster than $1/\sqrt{N}$. A modeler might consider variance reduction for computationally intensive simulations when practically attainable numerical accuracy dominates the conceptual or formalism uncertainty. For more discussion of all aspects of random number generators and their use in Monte Carlo simulations, we refer the reader to Press et al. (2002).

A separate source of numerical errors is *statistical procedures*, such as regressions or the computing of statistical characteristics of financial and economic data (standard deviations, correlations, etc.). While a quantitative financial analyst should be aware of these errors and know their impact on the model output’s accuracy, their discussion is outside the scope of this book. We refer the reader to any standard text on statistics or time series analysis, for example, Hamilton (1994).

How can one address the problem of model risk? The first step is to heed the injunctions about the intrinsic limitations of quantitative modeling, as already discussed in section 3.2: Not only is any model’s output just one of many equally valid ways of valuing an asset, but the very use of a quantitative model should be seen as complementary to other, non-quantitative, approaches. The second step is to study the *sensitivity* of the valuation to model details. For instance, one might investigate how the calculated value changes with changing the types of regressions used, the data frequency, the length of the historical window, the rate of the time decay, and other model parameters. One might also investigate a model’s stability with respect to the choice of factors. For instance, most stable models’ output should not change significantly if the S&P 500 index is replaced with the Russell 1000, or even Russell 3000, or if, less trivially, a Barclays Capital (formerly Lehman Brothers) fixed-income index is replaced with a Merrill Lynch one with similar characteristics (credit quality, geographic region, time to maturity, etc.). If there is a significant change, then it should be understood and explained in terms of the model’s specificities. In general, the magnitude of this change provides a measure of formalism uncertainty with respect to a particular exposure.

If Monte Carlo simulations are used, one should check that the doubling, say, of the number of trials leads to an economically immaterial change in the outputs that is within the expected statistical error range given by equation (3.1). A good model should be robust with respect to sensible parameter changes, that is, differences in the output for different parameters should be economically insignificant. More broadly, a family of sound models for the same problem should produce economically equivalent results; large differences may indicate a conceptual deficiency or implementational error.

In the remaining sections of this chapter we will make the preceding general discussion more concrete and quantitative before proceeding to apply these results to quantifying, measuring, and managing portfolio risks in Chapter 5.

3.3 VALUATION PRINCIPLES

The main purpose of quantitative financial modeling is twofold: to calculate assets' *intrinsic value* and to quantify this value's *response* to changes in external factors and asset parameters. Both conceptually and implementationally, valuation is the fundamental task: Having a ready valuation model, one can study the responses by varying the model's inputs.

Note that we are focusing here on complex assets whose value cannot be expressed in a closed form, that is, as a mathematical formula involving only elementary functions and, perhaps, other standard functions that can be easily and readily evaluated by most computer applications (e.g., the ubiquitous normal cumulative distribution function [CDF], also known as the "error function," defined in the Glossary, which features prominently throughout quantitative finance, in particular in the Black-Scholes formulae). When a closed-form solution exists, the sensitivities can also be expressed in a closed form. If available, closed-form solutions are always preferable, being more conceptually illuminating and computationally efficient; they are a rare example of objects that unite the normally antithetical qualities of being aesthetically appealing and utilitarian. Unfortunately, only a small number of idealized derivative instruments are amenable to exact analytical treatment, the best known example being the Black-Scholes pricing formulae for European options. Most derivatives and asset-backed securities need to be valued by DCF models of varying degree of complexity, often utilizing the Monte Carlo method.

A valuation model can be viewed as a "function of many variables": It takes multiple inputs and generates a single number, the asset's current value, a relationship that can be symbolically expressed as

$$V = \mathcal{M}(\mathcal{I}) \quad (3.2)$$

where V is the asset's current value according to the model, $\mathcal{M}$ represents the model's computational engine, and input $\mathcal{I}$ is the collection of all the data the model receives. We will distinguish between *continuous* and *discrete* factors. Furthermore, we differentiate between *market* and *non-market* continuous factors. Denoting the market continuous factors f_k ($k = 1, \dots, K$), the non-market continuous factors g_l ($l = 1, \dots, L$), and the discrete ones h_m ($m = 1, \dots, M$), we can write the input set as

$$\mathcal{I} = \{f_1, f_2, \dots, f_K; g_1, g_2, \dots, g_L; h_1, h_2, \dots, h_M\}. \quad (3.3)$$

The continuous factors are those that, like stock prices, market index levels, volatility, GDP, money supply, or unemployment level, are each expressed as a decimal or fractional number and can vary in increments much smaller than their typical values. The discrete inputs either are non-numeric ("liquid"—"illiquid," "law passed"—"law rejected," "buy"—"hold"—"sell," a credit rating) or can assume a small number of pre-determined—typically integer—numerical values (e.g., +2 for "strong buy," +1 for "buy," and so on until -2 for "strong sell"). When interested in current valuation one would use the latest available inputs; in order to observe response to input changes, one would shift these inputs in a controlled way.

In the remainder of this section we will describe the principles behind a typical model's computational engine $\mathcal{M}$. In this chapter's subsequent sections, we will show how to use an existing model to study an asset's responses to factor moves.

Discounted Cash Flows

The discounted cash flow (DCF) method of valuing assets assumes that an asset's economic value is the present value of the future cash flows the asset will deliver to its holder. The use of present value is an instance of the general principle of *time value of money*, a quantified statement of the fact that *almost always* (an exception, thus far unique in history, is described in the next paragraph) money today is better than the same money in the future. The reason is that if an investor had an amount of money $C(t_0)$ today (i.e., at time t_0), he could invest it, at least theoretically, in a zero-coupon government security maturing at a later time t and, with practically complete certainty, collect at that time a larger amount

$$C(t) = C(t_0) [1 + R_f(t, t_0)] \quad (3.4)$$

where $R_f(t, t_0)$ is the return on the government security (the “risk-free” return) over the holding period between the present, t_0 , and time t . Reversing this argument, if an investor expects to receive a sum of money $C(t)$ at a future time t , this sum’s value at an earlier time t_0 is a smaller number of units of currency

$$C(t_0) = \frac{C(t)}{1 + R_f(t, t_0)} < C(t) \quad (3.5)$$

because the investor could have made up the difference $C(t) - C(t_0)$ *risk-free* by investing in a zero-coupon government instrument (we do not discuss the tax consequences that can alter the argument in detail but not in substance; the standard reported Treasury spot rate curves typically account for the differences in tax treatment). That is, $C(t_0)$ is the *present value* to the investor, today, of a sum $C(t)$ to be received in the future, or, in other words, the investor is *indifferent* between receiving $C(t)$ at time t and $C(t_0)$ at an earlier time t_0 . The crux of the concept of time value of money is that, even though $C(t_0)$ is numerically smaller than $C(t)$, it is not less but *equally valuable* to the investor, precisely because it is received earlier: The difference $C(t) - C(t_0)$ is the time value, that is, the passage of time makes up for the difference in quantity. We will denote this relationship as

$$C(t_0) = PV_{t_0}(C(t)) \quad (3.6)$$

where the subscript t_0 indicates the “present time” to which the symbol $PV(\cdot)$ refers; we will often omit this subscript, when unambiguous, to improve readability.

The inequality in equation (3.5), as well as the whole argument of the preceding paragraph, presumes that risk-free rates, often proxied by yields on government securities, are always positive. In fact, on 10 days between December 4 and 26, 2008, one-month U.S. Treasury yields on securities trading in secondary markets became negative, while three-month yields went below zero on 9 days during the same period. The unprecedented occurrence of negative Treasury rates reflected the state of extreme fear that gripped the markets: Investors, not trusting even bank deposits and money market funds, were *paying* for the privilege of being a lender to the U.S. government, thus ensuring the preservation of most of their principal. This state of panic followed in the wake of the Lehman Brothers bankruptcy filing on September 15, 2008, which engendered fears of retail bank failures.

In the same month the shares of two established money market funds managed by Reserve Management Corporation “broke the buck,” or fell below \$1 in value.

Short-term U.S. Treasury yields turned negative again on November 19, 2009, although this was likely the result of financial companies’ scramble to burnish their balance sheet for the end of the year. The 2009 “window dressing” exercise, taking place in the ultra-low rates environment, was intensified by the fact that all leading U.S. financial companies reported their annual results at the same time at the end of December. Previously, some investment banks, including Morgan Stanley and Goldman Sachs, had their fiscal year end in November but had to change after becoming bank holding companies.

Negative sovereign yields, however, are an extreme event, unlikely to be repeated, and we will always assume them to be positive.⁷ Even the anomaly of December 2008 may, at least in part, have been a result of technical factors in the repo markets unrelated to time value of money. Note that the U.S. Treasury does not accept negative yield bids at U.S. security auctions.

Of course, an investor in a domestic government security, even in a developed economy, is still exposed to sovereign risk of default and to political risk, as discussed in section 2.2, as well as to inflation and reinvestment risks. A foreign investor would in addition be exposed to currency exchange risk. The concept of a “risk-free” investment is, nonetheless, useful, even if it means in reality only “the least risky,” or even “less risky.” Moreover, the *existence* of a risk-free rate at which market participants can borrow and lend without limit is an important, if imperfect, assumption of many derivative valuation models, including that of Black, Scholes, and Merton.

Generalizing to an asset expected to deliver N cash payments at times $t_1, t_2, \dots, t_N$, DCF states that this asset’s value at t_0 is the sum of present values of all the future payments:

$$V(t_0) = \sum_{i=1}^N PV(C(t_i)) \equiv \sum_{i=1}^N \frac{C(t_i)}{1 + R_f(t_i, t_0)} \quad (3.7)$$

⁷ In July 2009 Sveriges Riksbank, Sweden’s central bank (rikssbank.com), introduced a negative interest rate of -25 basis points on deposits kept with it by Swedish commercial banks. This unprecedented, and largely symbolic, measure, designed to encourage lending, does not imply negative nominal returns on Sweden’s government securities.

where, as before, $R_f(t_i, t_0)$ is the risk-free return over the holding period between t_0 and t_i .

It is customary and convenient to express risk-free returns in terms of annualized risk-free *rates* of return. We will use lowercase letters for annualized rates and, to avoid clutter, will assume *annual compounding*; for treatment of compounding with other periodicity, as well as for various day count convention, the reader can consult any standard text on fixed income securities, for instance Tuckman (2002). The relationship between rates and returns is given by

$$1 + R_f(t_i, t_0) = (1 + r_f(t_i, t_0))^{t_i - t_0} \quad (3.8)$$

where $r_f(t_i, t_0)$ is the risk-free annualized rate (with annual compounding) for the period between t_0 and t_i and time is measured in years. The expression Equation (3.7) for the DCF value of an asset can then be rewritten in terms of risk-free rates as

$$V(t_0) = \sum_{i=1}^N \frac{C(t_i)}{(1 + r_f(t_i, t_0))^{t_i - t_0}}. \quad (3.9)$$

Equation (3.9) encapsulates all the mechanics of DCF (compounding with non-annual periodicity would introduce somewhat cumbersome adjustments without changing anything conceptually). The non-trivial part of the DCF method is (1) choosing the discount rates $r_f(t_i, t_0)$ and, most importantly, (2) projecting the cash flows $C(t_i)$.

3.4 DISCOUNT RATES SELECTION

How can one obtain risk-free rates in practice? When valuing a financial asset, the analyst needs to discount cash flows occurring typically on a multitude of future dates (in fact, the payment dates themselves may change between different Monte Carlo trials). The appropriate discount rate for each cash flow $C(t_i)$ is determined by the return $R(t_i, t_0)$ on a *zero-coupon* risk-free security maturing exactly at time t_i . Of course, there does not exist a “risk-free” actively traded zero-coupon bond for each requisite maturity. The common practice, therefore, is to construct a continuous risk-free rates curve starting from a set of *reference rates*, namely, observable yields on liquid instruments considered to be to a reasonable approximation risk-free, and interpolating between the available maturities.

Reference Rates

In the United States it appears natural to use yields on Treasury securities as observable measures of risk-free rates. These yields, however, cannot be used directly since Treasury securities with tenors longer than one year pay coupon and, therefore, their yields, which effectively combine a series of shorter term rates, are not an adequate measure of zero-coupon rates. The solution is to construct a theoretical *spot rate curve*, that is, a curve of yields on hypothetical zero-coupon bonds, from active Treasuries markets. There are four commonly used approaches:

1. Use active (“on-the-run”) Treasuries—that is, use the set of maturities for which the Treasury issues bills, notes, or bonds, and for each of these maturities use the yield from the most recent auction.
2. Use on- and off-the-run Treasuries—that is, supplement on-the-run yields with additional points along the yield curve for which yields on outstanding issues with approximately the right time to maturity are used.
3. Use the Separate Trading of Registered Interest and Principal of Securities (STRIPS) market,⁸ where zero-coupon yields are in fact directly observable.
4. Use all outstanding Treasury securities, including STRIPS, to utilize all available market information.

Figure 3.3 illustrates the different curves commonly used in the U.S. markets.

Each approach has its advantages and disadvantages. For instance, while STRIPS may appear the most straightforward choice because they offer directly observable zero-coupon yields, these securities can be quite illiquid, that is, their yields reflect a liquidity premium, and quoted prices may diverge significantly from executable ones. In addition, the use of STRIPS is compounded by their tax treatment: The holder is required to pay tax on accrued interest even though he receives no cash, and this disadvantage depresses the prices. Finally, the picture is complicated even more by the fact that STRIPS issues created from the principal and those created from the coupons are subject to different tax treatment in some non-U.S. jurisdictions. Specifically, the gain on principal, but not coupon, STRIPS can be treated

⁸ While strictly speaking “STRIPS” is an acronym whose grammatical number is singular, signifying a way of trading components of Treasury notes and bonds, we will follow the established and convenient, if grammatically dubious, practice of using this word as a plural noun denoting the collection of securities traded in the STRIPS market.




FIGURE 3.3 U.S. Treasury yields, STRIPS yields, and LIBOR swap rates on October 8, 2009.

Source: Bloomberg Finance LP

as capital gain and taxed at a lower rate, leading to a yield bifurcation for the same maturity (for this reason, analysts often exclude principal STRIPS from spot-rate curve calculations).

Active, or on-the-run, Treasuries have good liquidity, and the issue selection for each maturity is unambiguous, namely, the one most recently auctioned. Moreover, Treasury bills (i.e., securities with tenor of up to one year) pay no coupon and therefore allow the direct observation of spot rates. Longer-term treasuries, however, do pay coupon, and their spot rates need to be calculated, in the simplest case by “bootstrapping,” a technique of calculating spot rates sequentially from a theoretical *par coupon curve*. The par coupon, or par yield, for a given bond is the estimated coupon rate that would make this bond trade at face value. The par rate is calculated for each Treasury issue from the observed yield; the more sophisticated methodologies adjust for the effect on the yield of the different tax treatment of interest and capital gain or loss—there is a capital gain on a bond purchased at a discount and a capital loss on a bond purchased at a premium (assuming they are held to maturity). As we see, the transition from observed active Treasury yields to spot rates requires two analytical procedures, both of which can involve non-purely mechanical steps: determining the tax treatment effect on the observed yields and calculating the spot rates, especially if techniques more sophisticated than bootstrapping are used; such steps are especially

prone to concept and formalism risks and therefore introduce additional model uncertainty. An arguably bigger problem, however, is that on-the-run tenors leave long gaps in the yield curve, most prominently between 10 and 30 years, that need to be interpolated.⁹ Figure 3.3 demonstrates the extended sections of the yield curve that lack active Treasury issues.

The long gap problem is often addressed by supplementing Treasury actives with *off-the-run* issues, that is, those whose tenor at issuance was not the desired point on the yield curve but whose remaining time to maturity is. The assumption is that if the Treasury auctioned a new security with the required tenor, it would sell at the same yield. There are two main problems with this procedure: first, there is ambiguity (and, hence, more model uncertainty, of the formalism approximation kind) in selecting the off-the-run security and, second, off-the-run issues often have lower liquidity and, therefore, distorted yields.

Some market practitioners believe that using only on-the-run and a few selected off-the-run Treasury securities wastes much of the information embedded in the complete Treasuries markets. Accordingly, there exist methodologies, extending beyond bootstrapping, for calculating the spot rate curve using data on all the outstanding Treasury securities. These methodologies adjust for tax effects and other special considerations, such as the effect of the active repo market on some maturities, and can incorporate the effect of *interest rate volatility* on the interest rate term structure, for example, by using stochastic modeling. While striving to reflect all the observable information, such techniques are of necessity complex and subject to the model risks we discussed in section 3.2.

In developed non-U.S. economies with liquid markets in local government securities, market participants have available to them sovereign spot rate curves constructed in a similar manner.

As a practical matter, major financial institutions often borrow directly from each other. This is particularly relevant to companies transacting in non-domestic currency, for example, banks in Europe transacting U.S. dollars. An established measure of the rates paid by borrowers in these transactions is the LIBOR. These rates are fixed every day shortly after 11 am London time by the British Bankers' Association (BBA). The BBA calculates LIBOR rates in 10 currencies for 15 maturities, ranging from overnight to 12 months, by polling a distinct panel of banks for each currency; the details of the methodology are described on the BBA's website (bbalibor.com).

⁹ At present, the Treasury issues bills maturing in 4, 13, 26, and 52 weeks; notes maturing in 2, 3, 5, 7, and 10 years; and bonds maturing in 30 years. Notes and bonds pay coupons.

For maturities longer than one year, the market perception of borrowing rates can be gleaned from LIBOR swap markets. The standard contract is on three-month LIBOR rates, and market quotations are available for a range of maturities extending to 30 years, as illustrated in Figure 3.3. Before they can be used for discounting, just as Treasury curves, the LIBOR swap rates need to be converted into spot (zero-coupon) rates by bootstrapping or a more sophisticated procedure.

The *LIBOR curve*, also known as *LIBOR swap curve*, combines the actual LIBOR rates, set once a day, by polling, before North American markets open, and continuously observable live market rates paid on the fixed leg of interest rate swaps. This combination of two disparate methodologies makes the two parts of the curve inconsistent (and, in particular, accounts for the pronounced shift in the LIBOR swap curve between one and two years, evident in Figure 3.3). This problem can be addressed by replacing LIBOR rates with rates deduced from short-term interest rate derivatives, such as Eurodollar futures, traded on the Chicago Mercantile Exchange (CME), or forward rate agreements (FRAs). In practice, many financial firms have interest rate desks that release “official” swap curves based on the prevailing market conditions, including in over-the-counter markets; these curves are then used to mark to market derivatives and other complex financial instruments.

Since investors’ borrowing costs are typically benchmarked off LIBOR, the use of LIBOR swap curves for asset valuation has been increasing despite the fact that, until recently, the swap curve market was often illiquid compared to sovereign debt (and, of course, LIBOR swaps are even less “risk-free”). As it happens, in many developed economies the swap market has become highly liquid, with tight bid-ask spreads, in some cases offering better liquidity than government bonds. In addition, the swap curve offers the following advantages:

- The lack of government regulation makes cross-country analysis easier. In particular, the market yields are not shifted by different tax treatment, as they are for sovereign debt.
- Since there is no underlying government bond, the rates are not affected by technical factors, such as the effects of the repo market. The rates are determined only by the market participants seeking to enter a swap at a given time.
- In contrast to government yields, LIBOR swap rates are not affected by sovereign credit risk. Instead, they reflect the credit risk of participating financial institutions, which is similar across countries, thus making cross-market comparisons easier.
- There are more quoted maturity points than available in active government debt markets, as illustrated in Figure 3.3.

Discount Curve Interpolation

Thus far we have discussed the determination of rates at a few discrete maturity points on the yield curve (the reference rates). As mentioned at the beginning of this sub-section, the valuation of realistic assets requires discount rates for a multitude of maturity points; in practice, one needs a *continuous* curve.

The commonly used interpolation methods range from linear (i.e., connecting the reference rates with straight lines, as in Figure 3.3) through spline fitting to fitting highly nonlinear functions, such as the Nelson-Siegel and Svensson templates. (A spline is a smooth, piecewise polynomials function.) In addition, stochastic methods of term structure determination use observable rates (usually all or most empirically available, with adjustments for tax treatment, liquidity premia, and technical considerations) as inputs to dynamic rate volatility models.

Each method has its advantages and disadvantages and may be better suited for a particular purpose. For instance, linear interpolation passes through all the empirical data points and is easy to implement but exhibits unrealistic kinks at observed maturities and, by its nature, is too simplistic to represent rates accurately between observations. On the other extreme, nonlinear models result in smooth continuous curves but may not pass through the empirical points. More worryingly, these models tend to smooth abrupt jumps in observed reference rates, thus hiding potentially valuable market information contained in yield curve irregularities. As a general rule, smoother curves, produced by more sophisticated fitting algorithms, are better suited for comparative analyses across markets or time horizons, where one wants to filter out the local and transient market aberrations, but not for marking to market, where the current market peculiarities and dislocations can make all the difference. A reasonable, and generally accepted, compromise is the use of cubic spline interpolation that offers satisfactorily smooth curves that pass through every empirical data point.¹⁰

Stochastic models of the interest rate term structure strive to analyze the future movements of interest rates, rather than infer, statically, the

¹⁰ Cubic splines' *second derivatives* (in the differential calculus sense) have kinks. This is usually not a concern but, if it is, then one can push the kinks to a higher—third—derivative by using quartic splines. In general, if one uses a spline of order n , then the derivative of order $n - 1$ would have kinks and the n th derivative would be discontinuous. One could always fit a single polynomial of order $N - 1$ to a curve with N reference rates; this polynomial will have all its derivatives smooth and continuous in the fitted region while passing through every observed data point. Unfortunately, high-order splines tend to exhibit wild oscillations between reference points that make them unsuitable for curve fitting.

appropriate discount rates from the current market. Consequently, it is not indisputable that these models produce intrinsically superior spot rates for the purpose of cash flow discounting in the context of asset valuation (as opposed to, say, making bets on interest rate moves or quantifying interest rate exposures): The concept of time value of money implies the ability to lock in the rates at the time of the valuation. When used, the stochastic, usually Monte Carlo, treatment of interest rates should ideally be consistent with the stochastic estimation of future cash flows (described in the next sub-section). For example, a Monte Carlo trial that exhibits a sustained decline in equity markets should, with high probability, also generate low short-term rates. In practice, implementing this inter-market coupling may be prohibitively expensive, in terms of both development effort and model running time.

For a more detailed discussion of term structure models the reader can consult, for instance, Tuckman (2002), Fabozzi and Fong (1994), and Fabozzi (2002).

3.5 CASH FLOW PROJECTION AND ASSET VALUATION

Projecting future cash flows is by far the more difficult of the two non-mechanical elements of the DCF approach. It requires a profound understanding of the asset, its interdependence with the broader markets and economy, with political and social developments, as well as the ability to translate these insights into the expected payouts and their timing.

The forthcoming cash flows from complex assets are determined by the future behavior of continuous and discrete factors introduced in equation (3.3), such as market prices, interest rates, macroeconomic indicators, political developments, credit events, and so on. The important point is that *given the future values* of all the relevant factors over the lifetime of an asset, the cash flows it generates can be calculated straightforwardly. The challenge, of course, is predicting the factors' future behavior. Absent a crystal ball or privileged access to a reliable oracle, the best an analyst can do is estimate the *probabilities* of the factors' assuming various values. With the probabilities available, the analyst can calculate future cash flows and their aggregate discounted present value for each future realization of the *factor universe* and then take the *average over the realizations*. The process of cash flow projection can therefore be divided into the following conceptual steps:

1. Identify the *factor universe* relevant to the asset. In addition to financial factors, it can also comprise political developments, legislative actions,

weather patterns, litigation outcomes, and others, determined by the nature of the asset.

2. Consider all the distinct *paths* the factor universe can take over the lifetime of the asset. A path is defined by a unique value assigned to every factor in the universe at each point in time during the asset's lifetime (in practice, usually on a discrete grid of points in time, e.g., daily or monthly).
3. Ascribe a *probability* to each path. This is usually the most difficult step. It requires discretionary input from the modeler, has considerable room for subjectivity, especially for discrete factors, and is particularly susceptible to the conceptual and formalism types of model risk.
4. For each path calculate the *cash flows* generated by the asset and their *present value*, equation (3.9), using the selected discount curve.
5. Calculate the *expected present value* as a weighted sum of present values for each path using the path probabilities as weightings.
6. Since the modeler has now at his disposal a set of valuations for all the paths (i.e., all the possible future realizations of the factor universe) and a set of corresponding probabilities, that is, the complete *probability distribution function* of asset values, he can calculate, at least in principle, any of the asset value's statistics: mean, median, standard deviation, skewness, kurtosis, quantiles, and so on.

Note that we treat the discount curve as given and determined extraneously to the cash flow projection procedure. As remarked at the end of the previous section, this approach may be not only practically easier but also conceptually preferable.

One of our assumptions is that calculating cash flows for each path is straightforward, even if the mechanics can become messy. What matters is that cash flows for a given path are *deterministic*, that is, for each path, for a given model, the set of cash flows is unique; there are no probabilities involved. Deterministic, though, does not mean free of uncertainty: The model itself is subject to uncertainties discussed in section 3.2 and, consequently, the cash flows are affected by model risks. The process of allocating cash receipts, ending in cash flows to the security holder after all other claims have been satisfied, is known as *cash flow waterfall*, a concept already introduced in section 2.2. For a discussion of waterfalls and the practical details of their implementation, the reader can consult Preinitz (2009).

Pre-Averaging Approximation

Occasionally analysts perform a *single-path calculation* in which only one typical, or representative, path is chosen for the factor universe. The sum of

the discounted cash flows for this particular path is taken as the estimate of the asset's value. Although it may appear that the single-path approach obviates the need for probabilistic analysis of the factor universe's alternative paths, this analysis is, in fact, embedded in the single-path method, explicitly or implicitly (and, sometimes, unbeknownst to the analyst). An analyst employing this method makes use of the *factor pre-averaging approximation* in the sense that he calculates (or guesses) the average value of each factor at every point in time and then inserts these averages into the algorithm for cash flow calculation.

The single-path method is usually easier to implement than the full stochastic "post-averaging." It can serve either as a useful "quick-and-dirty" first pass at a problem in preparation for the full-scale stochastic modeling or as an acceptable actual valuation approach when the factor volatility is low, high accuracy is not required, or resources are limited. One should never forget, however, the additional approximations inherent in this method. One disadvantage of a single-path calculation is that, unlike stochastic modeling, it does not automatically produce the probability distribution of asset values and, hence, their statistical characteristics.

Often models use pre-averaging on only some of the factors. As a simple example, consider a bullet corporate bond. A model that assumes, regardless of market and macroeconomic conditions, the same probability of absolute default (we disregard partial default for simplicity) and recovery rate for every month in the bond's life pre-averages these two quantities (estimated, most likely, literally as historical averages over bonds with similar characteristics). The pre-averaging in this case is only partial: The model still has distinct paths with distinct cash flow patterns, even though bond analysts usually do not think in these terms, each path corresponding to default occurring in a particular month, or not at all. We will make the concept of pre-averaging approximation more concrete in the next section.

3.6 STOCHASTIC ASSET VALUATION

When stating the mechanics of the discounted cash flow valuation in equations (3.7) and (3.9), we presumed that there was a unique estimate for the cash flow $C(t_i)$ at any given time t_i . We now assume that the future can proceed along many different paths, which we label with ω , corresponding to distinct realizations of the factor universe, each with a distinct sequence of cash flows¹¹ $\{C_\omega(t_i)\}_{i=1}^N$, and we assign a probability $p(\omega)$ for each of

¹¹ Cash flows can occur at different times in different realizations of the future, so that a formalist might write something like $\{C_\omega(t_i^{(\omega)})\}_{i=1}^{N^{(\omega)}}$. We eschew such pedantry

these alternative futures to occur. In other words, we admit that we cannot predict the future and set ourselves a more modest, albeit still daunting, task of describing the possible distinct paths the evolution of the factor universe can take and deciding on the likelihood of each.

The determination of probabilities $p(\omega)$ is in general a complex procedure, whose details depend on the specifics of the asset and on the valuation methodology. In fact, these probabilities need not even be the ones that the factors' future values can be reasonably expected to obey. The principle of *risk-neutral valuation*, widely used to price financial derivatives, states that the probabilities of future price moves should be those that would pertain in a hypothetical risk-neutral financial universe where investors do not require a premium for bearing risk. One implication is that the *expected return* on all assets, including the risky ones, should be set to the risk-free rate r_f . The investors' risk preference, expressed, for instance, as a required rate of return μ on a risky asset such that $\mu > r_f$, does not enter the discussion.

Risk-neutral valuation is so firmly established and its use is so ubiquitous that it is easy to forget that it relies on a number of assumptions. One is that the variations in the underlying security (or securities) and the derivative instrument are driven by the same stochastic processes. That is, to use a specific example, the source of randomness that causes fluctuations in a stock's price also causes fluctuations in the price of an option on this stock, *and there are no additional sources of randomness* behind the variations in the option's price.¹² If this is the case, then at any given moment, it is possible to construct a portfolio, consisting of the underlying security and the derivative, such that the stochastic innovations mutually cancel—this is a central point of the Black-Scholes-Merton framework for derivatives pricing. As the resulting portfolio has no stochastic component (i.e., is *risk free*), it must therefore return the risk-free rate to avoid creating an *arbitrage opportunity*, that is an opportunity for risk-free profit: If the risk-free portfolio delivered a return $r_{\text{portf}} > r_f$, then an investor could borrow at r_f , invest in the portfolio and book an assured profit of $r_{\text{portf}} - r_f$; if $r_{\text{portf}} < r_f$, then the investor could short the portfolio and lock in a profit of $r_f - r_{\text{portf}}$. The absence of arbitrage opportunities, or the markets' ability to respond nearly

and observe that if one elects the set of payout days $\{t_i\}_{i=1}^N$ to be all-inclusive, that is, to contain all cash flow dates in all possible realizations of the future, then one can simply set $C_\omega(t_i) = 0$ for those realizations ω without a payment at time t_i .

¹² This assumption often does not hold even for plain vanilla options traded on exchanges: Liquidity considerations and the market participants' ever-shifting desire to be exposed to specific expiration dates and strike prices, on the bullish (long call or short put) or the bearish (long put or short call) side, can lead to option price fluctuations not directly related to the price moves of the underlying stock or index.

instantly to eliminate them, is another assumption. The investor's ability to borrow and lend without limit at the same risk-free rate r_f is a third.

A consequence of the existence of the risk-free portfolio described above and of the arbitrage-free assumption is that the equation governing the evolution of this portfolio's price Π cannot contain the expected return μ on the underlying risky asset—this return, entering the separate equations for the underlying security's and the derivative's price evolution, should cancel alongside the stochastic innovations in the equation for Π . In fact, μ does not matter, and one can correctly price the derivative assuming any value for the risky asset's expected return, for example, r_f . This observation, which gives the principle of risk-neutral valuation its name, is inconsequential in practice when pricing through the use of Black-Scholes-Merton-type equations, which simply do not contain μ . The full practical power of risk-neutral valuation is revealed in pricing methods, such as multinomial trees and Monte Carlo, both described below, in which the calculation of cash flows (which do depend on μ) and their discounting are explicitly separated. The risk-neutral valuation principle states that one can assume the underlying instrument's expected rate of return to be r_f as long as the same r_f is used for discounting. One would arrive at the same valuation through, generally, a more complex process by using a "true" expected return $\mu > r_f$ and discounting with an appropriate *required rate of return* different from r_f (Hull, 2008).

The risk-free portfolio constructed of a derivative and its underlying instruments (a derivative can in general have multiple underlying securities) remains risk-free only for a short period of time: As the markets evolve, the portfolio composition needs to be continually adjusted, and the resulting trading activity is bound to incur transaction costs. These frictional losses are ignored in risk-neutral valuation.

While the assumptions behind risk-neutral valuation can be satisfied reasonably well in liquid markets for "traditional" financial derivatives, they often do not hold for complex asset-backed securities or portfolios of investments (which can be viewed as a kind of ABS): ABS markets frequently exhibit illiquidity and high transaction costs, while asset-backed securities are typically driven by multiple sources of randomness. In addition, discrete factors (e.g., the passing or rejecting of a law regulating the functioning of government-sponsored enterprises) and non-market continuous factors (e.g., GDP or unemployment rate) tend to play an important role in ABS valuation. The probability estimates for these factors' outcomes are determined by models usually not directly related to the underlying security or to the risk-free rate. The analyst should remember that risk-neutral valuation is itself a model that relies on a number of assumptions that may not hold, especially for complex financial instruments.

In practice, asset valuation models often do assume the expected rate of return on underlying risky assets, such as market indices, to be over the long term the risk-free rate, although this is more likely to be an expression of financial conservatism than an application of the risk-neutral valuation principle. Indeed, *expected* risky assets' returns must exceed the risk-free rate if they are to induce investors to put their money into anything other than government bonds. Empirically, the existence of a *risk premium*, that is an excess of the risky over the risk-free returns, has been confirmed in many markets over the long term. There are, however, notable exceptions. For instance, it is hard to preach the virtues of long-term investing to those who bought Japanese equities at the end of 1989 to find their investment, two decades later, worth less than a quarter of the purchase price.

The probabilities of discrete and non-market continuous factors, as used by typical ABS models, often depend on the modeler's perceptions or the details of macroeconomic models and tend to be subjective.

Now we can rewrite equations (3.7) and (3.9) for a specific path:

$$V_{\omega}(t_0) = \sum_{i=1}^N \text{PV}(C_{\omega}(t_i)) \equiv \sum_{i=1}^N \frac{C_{\omega}(t_i)}{(1 + r_f(t_i, t_0))^{t_i - t_0}} \quad (3.10)$$

where, to remind the reader, annual compounding is understood and time is measured in years. (We take the discount curve as given, i.e., not calculated stochastically via a coupled process described in section 3.4.) The resulting estimate $\bar{V}$ of the asset's worth is the expected value of its valuations, namely, the sum of V_{ω} s for all the realizations weighted with these realizations' probabilities:

$$\bar{V}(t_0) = \sum_{\omega} V_{\omega}(t_0) p(\omega) = \sum_{i=1}^N \frac{\langle C_{\omega}(t_i) \rangle_{\omega}}{(1 + r_f(t_i, t_0))^{t_i - t_0}} \equiv \langle V_{\omega}(t_0) \rangle_{\omega} \quad (3.11)$$

where we introduced the curly brackets notation for averaging over possible future factor universe paths ω : $\langle C_{\omega}(t_i) \rangle_{\omega} \equiv \sum_{\omega} C_{\omega}(t_i) p(\omega)$ and similarly for other stochastic (i.e., ω -dependent) variables. The probabilities $p(\omega)$ can be calculated explicitly or can be inherent in the stochastic valuation procedure.

Pre-Averaging Revisited

We now revisit the concept of pre-averaging introduced in section 3.5 in the context of single-path valuation. The *pre-averaging approximation* “collapses” the multiple realizations of the factor universe onto a single path.

The details of a given realization (path) of the factor universe determine a cash flow at time t_i through the factor values at all times up to t_i . That is, $C_\omega(t_i)$ is a function of $\mathcal{I}(t)$ for all $t \leq t_i$, where $\mathcal{I}$, defined in equation (3.3), is the set of all the factors in the factor universe. We can express this relationship symbolically as

$$C_\omega(t_i) = C(t_i; \{\mathcal{I}_\omega(t)|t \leq t_i\}),$$

$$\mathcal{I}_\omega(t) \equiv \{\{f_{k,\omega}(t)\}_{k=1}^K; \{g_{l,\omega}(t)\}_{l=1}^L; \{h_{m,\omega}(t)\}_{m=1}^M\}. \quad (3.12)$$

Pre-averaging consists in replacing the average cash flows

$$\langle C_\omega(t_i) \rangle_\omega \equiv \langle C_\omega(t_i; \{\mathcal{I}_\omega(t)|t \leq t_i\}) \rangle_\omega \quad (3.13)$$

in equation (3.11) with cash flows for the *average factors*:

$$\langle C(t_i; \{\mathcal{I}_\omega(t)|t \leq t_i\}) \rangle_\omega \approx C(t_i; \{\langle \mathcal{I}_\omega(t) \rangle_\omega | t \leq t_i\}). \quad (3.14)$$

Since cash flows usually depend on factors in a highly nonlinear way, this approximation is in general too crude to be useful. An important exception is when the variability of factors in the universe is small. Indeed, in the limiting case when for any given time each factor can assume only a single value, that is, $\mathcal{I}_\omega(t) = \mathcal{I}(t)$ for all ω and, therefore, there is no randomness, an equality is achieved in equation (3.14). Of course, the measure of “smallness” depends on the asset and the desired accuracy; in the case of discrete factors, one should consider the effect of the factor change, *ceteris paribus*, on the calculated cash flow rather than the magnitude of the change in the factor itself.

It may happen that the variability of only a subset of factors is small. In this case, a good approximation could be achieved with partial pre-averaging, as illustrated in section 3.5 with the example of a bullet corporate bond.

Mathematical Tools

What mathematical tools can one use to express and study the stochastic properties of asset prices? In quantitative finance it is usually one of the following:

- *Stochastic differential equations* (SDE), which explicitly incorporate random innovations.
- *Partial differential equations* (PDE), which directly relate price changes to changes in factors.

- *Multinomial* (usually *binomial* or *trinomial*) *trees*, a technique in which the possible future paths followed by the price of the underlying instruments, interest rates, or factors are confined to a predefined grid.
- The *Monte Carlo method* in which the asset is evaluated for a large number of sample realizations of the factor universe, drawn either from a preset distribution or from a pool of historical realizations.

The first two approaches are closely linked, and, in fact, for a broad and important class of problems, in particular when the price of underlying instruments obeys the log-normal distribution, there often exists a correspondence between SDEs and PDEs, so that the analyst can choose one or the other based on convenience, ease of formulation, and the analyst's preference. Differential equations govern the evolution of an asset's value $V(t)$ by relating small changes in V to the changing factors and to *stochastic innovations*; the latter enter SDEs explicitly and are implied by PDEs indirectly. The asset's current value $V(t_0)$ is obtained by solving these equations subject to the appropriate *boundary* and *initial conditions*.¹³ Stochastic and partial differential equations allow conveniently expressing the perceived or hypothesized dynamics governing prices or interest rates and are useful tools for conceptual investigation or as a way to set up a framework for a numerical solution. Unfortunately, while many quantitative finance models can be stated in terms of SDEs or PDEs, only a handful of exact or useful approximate solutions are known, examples being the Black-Scholes solution for European options (exact) and the Black approximation for American options, as summarized by Hull (2008). In addition, the stochastic or partial differential equation framework can be difficult to apply to complex, especially asset-backed, securities or to portfolios of investments: These assets are typically driven by a large number of stochastic innovations that are translated into the changes in $V(t)$ in an involved, circuitous way.

The multinomial tree approach, multinomial tree, pioneered by Cox, Ross, and Rubinstein (1979), is a powerful, versatile, and widely used tool for pricing financial derivatives, especially *contingent claims*, that is, options or other securities having the features of an option (e.g., callable or puttable bonds, convertible bonds, warrants, convertible preferred stock, interest rate floors and caps). The multinomial tree approach can yield a numerical solution to the problem of pricing *non-path dependent* contingent claims of

¹³ A good formal introduction to stochastic differential equations and their application in finance can be found in Øksendal (2000); for a general exposition of the use of SDEs and PDEs in finance, the reader can consult Joshi (2004) and Merton (1999).

almost arbitrary complexity, limited in principle only by the development cost and computing power.¹⁴ The absence of path dependence means that at every point in time, the value of the contingent claim depends only on the price of the underlying instruments at this time and not by these instruments' prior price history; examples of contingent claims without path dependence are European, American, and Bermudan options.

The multinomial tree method is not well suited for valuing path-dependent securities, such as Asian and look-back options, or those dependent on many stochastic variables (factors). Path-dependence could in principle be addressed through the use of so-called *non-recombinant* trees. These are much more complex than the commonly used *recombinant* ones, and the additional development effort and computing power that non-recombinant trees require usually render them impractical. Multiple stochastic variables would require numerous interdependent trees (roughly, one tree per each factor) with intricate interactions among them; again, the added complexity makes such models difficult to implement.

The value of complex asset-backed securities or investment portfolios usually depends on numerous factors and often is path-dependent (e.g., in the case of an MBS, the prepayment speed at a given point in time depends on previous interest rate history). A numerical valuation approach best suited for these securities is the Monte Carlo method.

3.7 THE MONTE CARLO METHOD

In this method, the asset is valued for a number N of randomly selected realizations $\Omega_N \equiv \{\omega_i\}_{i=1}^N$ of the factor universe drawn from the probability distribution $p(\omega)$. The method's name alludes to its similarity, at least superficial, with games of chance associated with the casinos of Monte Carlo, one of the communes of the Principality of Monaco. For each selected realization, or path, ω_i , we calculate the corresponding asset value and we use the average of these values as an approximation for the "correct" valuation, given in equation (3.11). That is the Monte Carlo method approximates equation (3.11) with

$$\tilde{V}(t_0) = \frac{1}{N} \sum_{i=1}^N V_i(t_0), \quad V_i(t_0) \equiv V_{\omega_i}(t_0). \quad (3.15)$$

¹⁴ A summary of the multinomial tree method and further references are available in Hull (2008). A useful introductory exposition of the underlying theory is given by Baxter and Rennie (2001).

TABLE 3.1 Normal Distribution Sampling Statistics

Statistics					Relative Errors			
N	100	1,000	10,000	Exact	N	100	1,000	10,000
Mean	5.300	5.020	5.010	5	Mean	6.00%	0.39%	0.19%
Median	5.141	5.088	5.016	5	Median	2.82%	1.77%	0.32%
Std Dev	2.177	2.037	2.008	2	Std Dev	8.86%	1.87%	0.42%
Skewness	0.124	0.015	0.023	0				
Kurtosis	-0.154	-0.183	-0.001	0	$N^{-1/2}$	10.0%	3.2%	1.00%

The *basic theorem of Monte Carlo integration* (Press, et al., 2002) implies that as N increases $\bar{V}$ will converge to the “exact” value $\bar{V}$ of equation (3.11) with the error decaying as $N^{-1/2}$, in accordance with equation (3.1). This error assessment is predicated on the assumption that the paths are chosen randomly from the factor universe distribution $p(\omega)$ or, equivalently, that the selections Ω_N are *representative* in the sense that the realized distribution $\tilde{p}_N(\omega_i)$ of ω_i s contained in each Ω_N is a reasonable approximation of $p(\omega)$ and that as N increases $\tilde{p}_N(\omega_i)$ converges to $p(\omega)$.

To illustrate this point, we drew samples from a one-dimensional normal distribution with mean 5 and standard deviation 3 (this simple situation of the factor universe comprising a single variable arises, e.g., in valuing a plain vanilla stock option when the stock volatility and risk-free interest rate are assumed to remain constant during the option’s lifetime). Figure 3.4 demonstrates a sequence of distributions realized for random samples¹⁵ for various values of N . This sequence appears to approach (converge to) the actual probability distribution function (PDF) as N increases. Table 3.1 shows some statistics calculated with the same samplings and these statistics’ scaled deviations from their exact values (the table does not list the relative errors for the skewness and the kurtosis since these statistics equal zero for a normal distribution). Note that the magnitude of the errors declines with increasing N , in line with equation (3.1). This observation concurs with the visible convergence of the measured PDF to the exact one.

Figure 3.4, however, disguises an important limitation of Monte Carlo simulations: their restricted ability to probe the *tails* of a distribution. Indeed, the four outer points of the $N = 100$ PDF are exactly zero, that is, not a single sampling point fell into these bins (or the regions further in the

¹⁵ In this special case of a single factor variable and a single time period, each path ω_i degenerates into a single number drawn from the specified distribution.




FIGURE 3.4 Normal probability distribution function with mean 5 and standard deviation 3 compared to realized distributions of random samplings of 100, 1,000, and 10,000 variables. The realized distributions were calculated as the fraction of randomly drawn numbers that fell in a bin of width unity centered on the indicated abscissa. For example, the square at 6 indicates that out of 100 random variables drawn from the normal distribution 12% fell between 5.5 and 6.5.

tails). Since the normal distribution extends indefinitely in both directions and at any point its value is non-zero, even if small, the sampled distribution underestimates the correct one in these outlying areas by a factor of infinity! Increasing the sample size would widen the probed range of variables, but for a given N there are always limits x_{Left} and x_{Right} , determined respectively by $NF(x_{\text{Left}}) \approx 1$ and $N(1 - F(x_{\text{Right}})) \approx 1$, $F(x)$, being the cumulative distribution function, on how far into the tails a sampling can extend.

We would like to emphasize that the limitation we have just described is a failure to incorporate the full details of the distribution underlying the Monte Carlo simulation, whether explicitly assumed or implied. It does not address the question of this distribution's adequacy for describing actual factor returns. It is instructive to compare this with the case of differential equations. In that case, a solution incorporates the behavior of the underlying distribution in its whole domain, including in the tails. But, just as in the

case of Monte Carlo, this underlying distribution may not provide a good description of realized market behavior.

But why should one worry about these outlying regions? After all, if the true probability of finding a variable there is very small, should they not be irrelevant? It turns out that the tails are important in calculating higher *cumulants* (also known as *semi-invariants*, see Appendix A for definitions): the skewness, the kurtosis, and even the standard deviation—note that in Table 3.1 the relative errors of the means (first cumulants) are smaller than the relative errors of the standard deviations (square roots of second cumulants) because the second cumulants probe further into the tail, where the statistics are sparser. The reason should be clear from the definition of higher cumulants given in Appendix B: Near the peak of the distribution function the integrand vanishes, while attaining maximum magnitude at points that recede into the tails as the order of the cumulant increases (the cumulant order equals 1, 2, 3, and 4 for the mean, the variance, the skewness, and the kurtosis respectively). But Monte Carlo's failings in exploring the outer reaches of probability distributions have particularly dire consequences for *tail risk measures*, as we discuss in section 5.4, especially the much overused and often insufficiently understood value at risk.

Given a sample of path valuations $\{V_i\}_{i=0}^N$, one can calculate approximate values of various statistics of V . For example, sample estimates of V 's variance, skewness, and kurtosis are calculated, respectively, as

$$\begin{aligned}\tilde{\sigma}_V^2 &= \frac{1}{N-1} \sum_{i=1}^N (V_i - \tilde{V})^2, \\ \tilde{\gamma}_V &= \frac{1}{N\tilde{\sigma}_V^3} \sum_{i=1}^N (V_i - \tilde{V})^3, \\ \tilde{\kappa}_V &= \frac{1}{N\tilde{\sigma}_V^4} \sum_{i=1}^N (V_i - \tilde{V})^4 - 3\end{aligned}\tag{3.16}$$

where we suppressed the time argument to reduce clutter and used the conventional non-dimensional definitions of skewness and kurtosis, wherefore the $\tilde{\sigma}_V^3$ and the $\tilde{\sigma}_V^4$ in the denominator. As an aside, we note that skewness and kurtosis vanish for a normal distribution and that a combination of negative skewness and positive kurtosis is potentially lethal as it signifies a high probability of large negative returns.

While a Monte Carlo calculation can in principle simulate both non-market continuous and discrete factors, the true power of the method lies

in modeling market variables: stock, bond, and commodity prices; equity and fixed income indices; interest rates (not to be confused with discount rates); foreign exchange rates. A typical model would sample market factors f_k while keeping discrete, h_m , and non-market continuous, g_l , factors fixed. In order to gauge the asset value's sensitivity to non-market factors, such as investor behavior, unemployment level, credit rating, or the passage of a law, the modeler would usually perform a number of calculations for a judiciously selected range of these inputs.

Types of Monte Carlo Simulation

Monte Carlo calculations can be classified into two main types with respect to the method they use to generate random market factors: *historical* and *analytical*. A historical calculation draws factor returns at random from a pool of previously observed ones, while an analytical Monte Carlo generates factor returns governed by a preset probability distribution whose parameters (means, volatilities, correlations, etc.) are either obtained historically or, in the case of implied volatilities and, less often, implied correlations, can be observed in the markets. Another, *hybrid* type, which can be considered a sub-type of historical Monte Carlo, generates random numbers directly from an empirically observed distribution of returns.

Historical Monte Carlo The principal appeal of the historical approach is that it uses factor returns that have actually occurred. Moreover, in multifactor models, at any given time step in the simulation, this method will select all the factor returns from the same historical time period, thus reproducing cross-sectional factor co-movement. The selected returns are put back into the pool, so it does not shrink as the simulation progresses.

Historical Monte Carlo, however, has its limitations. One of the most important ones is that it is ignorant of what has *not* occurred: It can only warn of the devil you already know. Another significant shortcoming is that, while preserving cross-sectional correlations, it loses the chronological sequence of historical market moves. This can distort the patterns of price evolution and, consequently, skew the value of path-dependent assets.

The modeler needs to decide on the historical period used as the source of the returns in the pool. It usually makes sense to end this period with the most recent observation but there is no clear-cut rule for how far into the past it should extend: Too short a period will provide limited statistics and may miss critical turbulent episodes, while reaching too far into history may sample periods when different market and economic conditions prevailed, as well as run into problems of data inconsistency—the company of interest

may not have existed or may have pursued a different line of business, the index composition may have changed, new regulations may have been enacted, and so on.

A related concern is that the compiled pool of factor returns can be too small to achieve the desired numerical accuracy of Monte Carlo simulations. For example, a pool of monthly index returns compiled from 40 years of data would contain 480 data points. This means that the accuracy of the calculation will stop improving after about 1,000 trials, that is, by the time each data point has been drawn on average once or twice. The reason for this is that the $N^{-1/2}$ rule of equation (3.1) applies for *independent random* variables. A pool of N_{pool} historical returns can be considered approximately independent and random, but a sequence of mN_{pool} returns drawn from this pool will still contain only N_{pool} independent numbers. For a large m , each member of the pool will be selected about m times, that is, such a simulation is equivalent to performing m identical Monte Carlo runs, each with exactly the same sequence of N_{pool} returns and producing the same valuation $\tilde{V}_{N_{\text{pool}}}$. No matter how big an m one chooses, the ultimate calculated value is the same $\tilde{V}_{mN_{\text{pool}}} \approx (m\tilde{V}_{N_{\text{pool}}})/m = \tilde{V}_{N_{\text{pool}}}$. Effectively, the N in equation (3.1) saturates at a small multiple of N_{pool} . (The presence of path-dependence would somewhat complicate the argument without changing its essence.) This numerical accuracy saturation may not, however, be a significant practical limitation since the numerical error achievable with a pool of modest size would often be subsumed by other model uncertainties.

Generating random numbers according to the *historically realized probability* is relatively straightforward for a single market factor f_k : One can easily construct an empirical cumulative distribution function $F(f_k)$ and then “invert” it in order to transform uniformly distributed random numbers into random variables governed by $F(\cdot)$ (Press et al., 2002). This hybrid approach overcomes the problem of small pool size that afflicts the standard historical Monte Carlo while retaining the benefit of the direct use of market observations. Inverting empirical probability for multiple variables, however, can become complicated, with the degree of complexity exponentiating as the number of modeled factors increases, and, as a result, hybrid Monte Carlo is rarely used for more than one market factor.

The average rate of return realized in a historical Monte Carlo calculation is determined by the sequence drawn from the pool, or by the empirical cumulative distribution function (CDF) in the hybrid case—it cannot be pre-determined and, therefore, this method is not suitable for risk-neutral valuation.

Historical Monte Carlo simulations should not be confused with *historical tests* introduced in section 3.2, in which historically realized data are used

as *inputs* into a model (which may, by chance, use historical Monte Carlo simulations, or may not use Monte Carlo of any kind) and the model's outputs are compared with subsequent historical observations. In that section, we discussed using historical tests as part of the model validation process. In the context of asset valuation, historical tests are a helpful tool for assessing an asset's performance under actual market conditions, especially during periods of financial stress.

Analytical Monte Carlo In *analytical Monte Carlo*, the modeler decides up front on the factor return probability distribution. The decision process has two parts: (1) the choice of the *functional form* of the probability distribution function and (2) the choice of the *parameters* that determine the location of the peak of the distribution, its width, the degree of asymmetry, the heaviness of the tails, and so on. Note that we avoided the names of standard statistical measures of these properties, namely, the mean, the variance, the skewness and the kurtosis. This was not by accident: There is evidence (Mantegna and Stanley, 2007; Voit, 2005) that the behavior of market returns can be better described by a family of probability distributions with tails so heavy that the variance and other standard statistics do not exist (or, more accurately, are infinite). These distributions bear the name of Paul Pierre Lévy, a French mathematician who described and studied them, and phenomena governed by them are known as *Lévy flights* or *Lévy processes*. While Lévy's student Mandelbrot (1963) observed almost five decades ago that logarithmic (i.e., continuously compounded) daily and monthly returns¹⁶ on cotton prices followed Lévy distributions, a physical system where Lévy processes arise naturally from the underlying dynamics was predicted theoretically by Bychuk and O'Shaughnessy (1995) and confirmed experimentally by Stapf, Kimmich, and Seitter (1995) much more recently. In today's liquid markets, Lévy distributions are evidenced to a good approximation on time scales up to minutes, hours, or days (Mantegna and Stanley, 2007).

Most Monte Carlo models, however, assume normal distributions for logarithmic returns (this implies that the prices are governed by log-normal distributions; see Appendix A for definitions). There are theoretical, practical, and empirical justifications for this. From the probability theory point of view, *if* price changes over short periods were identically distributed random numbers, independent, and *had a finite variance*, then logarithmic returns over longer periods would obey a normal distribution regardless of

¹⁶ If $S(t)$ denotes an asset's price at time t , then this asset's logarithmic price return between times t and $t + \Delta t$ is $\ln [S(t + \Delta t)/S(t)]$. It approaches the arithmetic return $[S(t + \Delta t) - S(t)]/S(t)$ in the limit of small Δt .

the distribution governing the short period changes; this is the essence of the central limit theorem (CLT), which we discuss in Appendix A. The special property of logarithmic returns that makes the CLT applicable is that unlike, say, arithmetic returns they are *additive* in time, as is clear from the definition in footnote 16: $r(t_2, t_1) = r(t_2, t') + r(t', t_1)$ where $r(t_2, t_1)$ denotes logarithmic return between times t_1 and t_2 , and $t_1 < t' < t_2$. It is this property that enables one to express longer-period logarithmic returns as sums of many shorter-period ones.

There are problems with all three assumptions needed for the normality of logarithmic returns. The classical formulation of the CLT requires short-term returns to be *identically distributed*, that is, drawn from one unchanging underlying distribution. In fact, as we know, capital markets have complex and evolving dynamics, and asset prices are certainly not drawn from one immutable PDF. One piece of evidence is the undeniable existence, on all time scales, of sustained and alternating bear and bull runs. The same evidence invalidates the second assumption: There is well-established momentum in market moves, that is, successive price changes are not independent. There exist extensions of the classical CLT that relax the identical distribution requirement (Shiryaev, 1995) and even, in exchange for small deviations from exact normality, allow a small degree of interdependence (Ma, 1985). These extensions, however, assume certain constraints on the relative magnitude of short-period returns and their degree of interdependence; these constraints are likely to be violated during abrupt market moves. Finally, as noted previously in this subsection, the variance of small moves may not exist in a meaningful sense.¹⁷ An additional complication is that under the CLT the convergence of the PDF of a sum of M random variables to the normal one occurs mostly in the bulge of the distribution, spreading into the tails relatively slowly with increasing M (the exact speed of convergence depends on the true probability governing the random variables).¹⁸ This contributes to the perennial “tail problem” in asset valuation.

¹⁷ If one assumes that a Lévy distribution governs price changes at small time scales, then there must be ultimately a “cutoff” for the heavy tails: There is, after all, a finite amount of funds available in the markets. This cutoff, however, would be vastly greater than the typical price moves, as would be the (finite) standard deviation, comparable with the size of the cutoff. This contrasts with the case of “ordinary” (e.g., normal) distributions where the standard deviation is of the size of typical moves (in fact, determines them).

¹⁸ This statement does not contradict the Berry-Esséen theorem, which asserts that, under certain constraints on random variables’ third moments, the *cumulative* distribution function of the sum converges uniformly to the normal CDF.

The practical ease of using normal distributions is probably the pre-eminent reason for their widespread deployment: Normal random variables, including correlated vectors thereof, are easy to implement and computationally inexpensive to generate. There exist standard and efficient algorithms for producing multidimensional normal random numbers with a given variance-covariance matrix; we outline these procedures in Appendix A. The main reason for this ease of use is that a normal distribution, for a single variable, is completely defined by only two parameters, the mean and the variance. The skewness, the kurtosis, and an *infinite number* of higher, unnamed cumulants all equal zero. Of course, one can construct any number of probability distribution functions defined by a small number of parameters, even if all the cumulants are non-vanishing (and finite). The complication is that now one needs to ensure that all the cumulants match the empirically observed (or market-implied) ones. In reality, one would match a few lower cumulants, probably up to order four (kurtosis), and ensure that the higher ones are not unacceptably large. Similarly, a joint normal distribution for multiple variables is determined, in addition to the means and the variances, only by bilinear pair-wise covariances. In other words, a multivariate normal distribution for K variables is completely defined by a vector of K means and a $K \times K$ symmetric variance-covariance matrix, containing $K(K + 1)/2$ independent parameters. By contrast, a general multivariate probability distribution is specified, in addition to the infinite number of cumulants for each variables, by an infinite number of covariance measures, combining various powers (up to infinity) of all the possible combinations of all the K variables. One more convenience argument in support of using the normal distribution is that its PDF has the form of a function, the Gaussian, that possesses helpful properties, including extremely fast decay beyond a few standard deviations from the maximum, which make inverting the normal CDF, for either a single or multiple factors, uniquely easy.

Our discussion of Lévy distributions and processes notwithstanding, logarithmic returns observed in liquid markets over longer time periods (days, weeks, months) and *under normal market conditions* to a good approximation do follow a normal distribution. This fact furnishes an empirical justification for the normal assumption underlying many valuation models in existence: It is acceptable to use normal distribution for many applications, provided that the user remains alert to its limitations, especially when markets behave *abnormally*. Furthermore, because of normal distribution's propensity to underestimate tail risk, its use is more justifiable for the purpose of routine asset valuation than for risk management, especially for calculating tail risk measures discussed in section 5.4.

While, as we have observed, a normal distribution's tails nominally extend to infinity, they decay exceedingly fast beyond a few standard deviations

(e.g., only 0.27% of the probability lies outside three standard deviations of the mean) and cannot therefore model unusually abrupt market moves. This property of the Gaussian gives rise to the occasional sensationalist stories in the media about “multiple-sigma events” and astronomically long waiting times required for such events to occur “in theory.” The “theory” in this case assumes a log-normal price distribution, and the extreme event simply refutes this assumption. The theoretical chances of these extreme moves occurring would agree better with empirical observations under the assumption that market prices obey the laws of Lévy statistics; unfortunately, the long-tail Lévy distributions are more difficult to handle analytically.

One might be tempted to use a non-normal distribution with a finite variance and non-vanishing finite skewness and kurtosis as a means of simulating tail events. This exercise would likely be futile: An important, and somewhat counterintuitive, consequence of the central limit theorem is that the distribution of returns over longer periods “forgets” the details of short-term return statistics; no matter how skewed or heavy-tailed the short-term distribution, as long as it has a finite variance (and other conditions of the CLT are satisfied), the distribution of longer-period returns becomes symmetric and light-tailed—as light as the normal distribution. By contrast, *Lévy limit distributions* (i.e., those, to which distributions of sums of Lévy variables converge—the normal is the limit distribution of variables with finite variance) do possess both heavy tails and, in general, asymmetry.

The normal distribution’s “short-tail” problem is exacerbated in the case of Monte Carlo, as discussed earlier in this subsection, by this method’s intrinsic inability to probe adequately even these diminutive tails. In short, a valuation method based on the normal distribution is liable to miss big price jumps, especially in turbulent markets when the methods used to calculate the distribution parameters are unlikely to catch up with the fast changing conditions.

Parameter Selection This leads us to the second aspect of analytical Monte Carlo modeling, namely, the task of choosing the parameters that, given the functional form, define the details of the underlying probability distribution. As mentioned, a model that assumes a normal distribution for K market factor logarithmic returns requires specifying K means and $K(K + 1)/2$ elements of the variance-covariance matrix. In the discussion preceding equation (3.10), we explained why most valuation models assume the expected rate of return to be the risk-free rate r_f . We will only note here that r_f should be the *total* rate of return on a risky asset. This means, for instance, that the price rate of return on a stock or stock index should be r_f reduced by the expected dividend yield over the lifetime of the asset.

The bulk of the task of parameter determination falls on calculating the variance-covariance matrix, that is, K variances and $K(K - 1)/2$ covariances. In practice, the model inputs are usually stated in terms of the more intuitive standard deviations (square roots of variances) and *correlation coefficients*, but they need to be converted into the variance-covariance matrix at some stage in the random number generation—see Appendix A for the conversion formulae and discussion of normal random number generation. Most of these $K(K + 1)/2$ model input parameters are calculated from historical time series, and the modeler faces problems reminiscent of those arising when selecting a pool of realized returns for historical Monte Carlo: He needs to decide on the input data frequency (daily, weekly, monthly, quarterly, or, in some cases, intraday), the length of the time series (the *valuation period*), and, a novel dilemma, the *weighting* of the historical data points. In a simplest, equally weighted, calculation all the points in the time series are given the same significance. This means that events that occurred years ago have the same impact on model inputs as those of yesterday—a model like this cannot react nimbly to changed circumstances. Another undesirable feature is that unusually big moves would abruptly change the inputs not only immediately after their occurrence, but also when they “drop out” of the valuation period, possibly many years after the event. There is, however, the advantage of using all historical data, at least within the valuation period, even-handedly, avoiding the modeler’s subjective preconceptions about the different data points’ relative importance.

A common alternative is to use *historical decay*. In this method, older data points are given less weight in the calculation. Specifically, their weighting is usually set to decay exponentially with age, the rate of decay characterized by a *half-life*. If, for example, the half-life is set to one year, then a return observed one year ago carries half the weight in the calculation compared to a return observed yesterday, and a return from two years into the past, one-quarter the weight. Historical decay puts more emphasis on recent events, while the influence of distant history gracefully fades away. The problem, of course, is that the older events may be exactly the ones to be concerned with: A precipitous market drop would exert practically no influence on the model inputs after several of half-lives. In addition, historical decay introduces into the model a new parameter, the half-life, whose determination relies on the modeler’s subjective judgment. (To be fair, one should note that selecting a half-life is not dissimilar to choosing the length of the valuation period.)

As is so often the case in asset modeling, there is no single right way to calculate the inputs; the modeler needs to use his judgment, intuition,

and experience to determine the type of calculation, the data frequency, and the valuation period. The characteristics of the valued asset and the purpose of the modeling exercise should provide direction. As a practical guide, the length of the valuation period (or the half-life of historical decay) should be commensurate with the asset's expected lifetime, while the frequency should be related to the expected frequency with which the asset will be valued or its risk measures calculated. At the model implementation stage, the developer should conduct historical tests for a range of input calculation methods and identify the best performing ones. Ideally, the model in production should then be run for a number of input evaluating methodologies, with the results providing complementary views of the asset's value and risks, some emphasizing the short-term pitfalls and others providing a longer term perspective.

Historically determined input parameters reflect the average price behavior over an extended period of time and may not respond to the latest events that are almost certain to change future market characteristics. For stocks and indices with active options markets, there may be a way to gauge the current market perception of future price variability through *implied volatility*. Since implied volatility is *observable*, it is an instantaneous measure of market sentiment and can be a better indicator of future price behavior than a historically calculated standard deviation. It is common, therefore, for models to use, where available, implied volatilities in lieu of historical ones for market factors such as actively traded indices. One should be aware, however, of two caveats. First, implied volatilities are *model-dependent*, that is, their value depends on the model used to value options. The modeler needs to ensure that, inasmuch as possible, his asset model is consistent with the option model. For example, the two models should use consistent dividend yield estimation methodologies. Second, observed implied volatilities use option prices, which reflect factors other than the underlying instrument's observable or perceived characteristics. For instance, these prices are affected by option liquidity considerations and by the market participants' flitting desire to be exposed to particular expiration dates and strike prices and to take bullish or bearish positions. Perhaps even more importantly, one should not forget that implied volatilities are first and foremost expressions of *market sentiment*, that is, the collective view of a group of human beings and computer programs, created by human beings, operating under time pressure and lacking perfect foresight.

With the process for calculating inputs established, the modeler usually updates them at the shortest intervals allowed by the calculation methodology and data availability. Thus, implied volatilities would typically be obtained before every run, while historically calculated variances

and correlations would be recalculated when the next datum becomes available (at the close of the day, the end of the week, month end, etc.).

For models using non-normal distributions with finite variance, the process of calculating the means and the variance-covariance matrix is similar to the one just outlined. The treatment of higher cumulants and covariances depends on the details of the asset and the complexity of the model. Typically, the model would use the skewness and the kurtosis of each factor, which are calculated historically, while ensuring that the magnitudes of other higher moments are contained.

Since Lévy distributions do not possess a finite variance (nor, often, a finite mean), models of Lévy processes require special treatment that is beyond the scope of this book.

Jump-Diffusion Process To conclude this subsection, we mention an alternative to Lévy distributions as a means to model large price moves: the *jump-diffusion process*, which combines normal (Gaussian) price diffusion with abrupt discontinuous jumps whose frequency is governed by the Poisson distribution, defined in Appendix A, while the magnitude of the jump can be itself a random variable (Joshi, 2004). In contrast to Lévy probabilities, the Poisson distribution is a fast-decaying discrete function of the number of jumps occurring during a given time period; its variance and all higher cumulants are finite. However, the introduction of a discontinuous process upsets some of the basic premises of the Black-Scholes-Merton valuation framework. In particular, the expected rate of return, that is, the investor risk preference, now enters the discussion, and, as a consequence, even for plain vanilla options the arbitrage-free argument yields not a specific price but rather bounds on possible prices; the actual price depends on the degree of risk aversion (or risk seeking) prevalent in the market.

3.8 STOCHASTIC EXTRAPOLATION

Inherent in all the stochastic valuation methods we have discussed in this section is the principle of *stochastic extrapolation*. Its basic premise is that while we cannot predict the exact evolution of the asset under consideration, or of the factor universe determining its behavior, we can make a meaningful statement about the expected *statistical properties* of the future states of this factor universe. Many valuation algorithms assume that these statistical properties, as determined by recent history, will remain unchanged over the life of the asset; more sophisticated treatments model changes in statistical properties themselves, for instance, by making volatility a stochastic quantity

or by conditioning volatility on returns. But in all cases, the idea is to distill historical and current market observations into a small number of statistical characteristics and project them into the future, either unchanged or evolving according to pre-specified rules. The reduction in the number of extrapolated quantities improves the accuracy of the factor universe projecting and, consequently, asset valuation.

CHAPTER 4

Market Exposures and Factor Sensitivities

... there is not a single effect in nature, even the least that exists, such that the most ingenious theorists can arrive at a complete understanding of it. This vain presumption of understanding everything can have no other basis than never understanding anything.

Galileo Galilei, *Dialogue Concerning
the Two Chief World Systems*

4.1 FROM VALUATION TO RESPONSES AND SENSITIVITIES

A tested and audited valuation model serves as a foundation for the second element of the *risk management triad* of identifying, quantifying, and managing market exposures that we set out to elucidate in the Preface. In this section we will show how to use a model to tabulate an asset's responses to large (or, more generally, *finite*) factor changes, as well as how to calculate the asset value's *sensitivities* to small (in the limit, *infinitesimally* small) shifts in continuous, especially market, factors and how to translate these sensitivities into *equivalent variation investments* (EVIs) that can be directly used for sizing hedges.

In section 3.3 we introduced the view of the valuation model as a computational engine $\mathcal{M}(\cdot)$ that processes inputs $\mathcal{I}$ to generate a value V , a relationship formalized by equation (3.2). When interested in current valuation, an analyst would use the latest available inputs $\mathcal{I}_{\text{current}}$ (or, if valuing an asset on a past date, the inputs observed at the corresponding time). In order to calculate *responses*, a risk manager would feed $\mathcal{M}(\cdot)$ a range of

input sets $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_Q$ distinct from $\mathcal{I}_{\text{current}}$ and examine the resulting valuations. The choice of *scenarios* $\{\mathcal{I}_q\}_{q=1}^Q$ depends on the desired type of response measures and the objective of the exercise, as we detail in the rest of this section.

While the risk manager should understand the principles underlying the valuation model and the model's limitations, for the purpose of calculating an asset's responses to factor changes he can treat the model as a "black box," that is, as an application that, given a set of inputs, produces a unique valuation, the details of this application's implementation being inessential.

Responses to changes in discrete factors $\{h_m\}_{m=1}^M$ can usually be evaluated straightforwardly by enumeration: Even for assets with multiple discrete factors, the number of realistic combinations of values of h_m is rarely prohibitive. Our main focus in the remainder of this chapter will be on *market factors* (responses to non-market continuous factors can be obtained similarly). To simplify notation we will omit non-market factor arguments in V and introduce vector notation for market factors:

$$V \equiv \mathcal{M}(\mathcal{I}) = V(f) \quad \text{where} \quad f \equiv \begin{Bmatrix} f_1 \\ f_2 \\ \vdots \\ f_K \end{Bmatrix}. \quad (4.1)$$

A risk manager, interested in knowing how market moves will affect the value of a portfolio, can approach the problem from one of two different directions. He can evaluate the portfolio for a range of market factor levels, in other words, create a *response matrix* or, alternatively, he can calculate the portfolio's *sensitivities* to market factors—a set of sensitivities enables one to evaluate the portfolio's response to arbitrary but small changes in vector f .

The response matrix (in general, multidimensional) has the advantage of presenting "exact" model valuations for arbitrarily large shocks to the markets. The main disadvantage is that a manager, looking at a response matrix, sees a frozen snapshot of asset values for a limited discrete set of market conditions. A collection of sensitivities, by contrast, allows the manager instantly to revalue the asset for any arbitrary combination of moves in the market factors. These new valuations, however, would be *approximations* to model values, with the error growing as the magnitude of market factor moves increases.

In the rest of this chapter we illustrate these points and elaborate on them.

4.2 RESPONSE MATRIX AND SCENARIO GRID

Consider an asset with U.S. and euro-zone equity exposures and let this asset's market factors be, accordingly, the levels of the S&P 500 index (f_1) and the Dow Jones Euro STOXX 50 index (f_2). An asset with such exposures could be, for instance, backed by mutual fund fees, or it could be a portfolio of options and other derivatives. Its value would, of course, depend on other market factors (risk-free rates as a minimum), which we ignore for simplicity. In addition, we do not explicitly list the volatilities of the two indices as distinct market factors, but we will discuss the effect of volatility changes on valuation and demonstrate a way of incorporating them in a response matrix. We assume that the asset manager is U.S.-based and that the Euro STOXX 50 is converted to U.S. dollars.

Response Matrix

The latest index levels¹ $f_{\text{current}} \equiv (f_1^{(c)}, f_2^{(c)})$ will produce the current asset valuation $V_{\text{current}} = V(f_{\text{current}})$. Suppose $V_{\text{current}} = \$100$ million. The asset manager (and, especially, his supervising risk manager) might reasonably wish to know how the value of the asset would change if the level of the two equity indices changed by various amounts relative to $f_1^{(c)}$ and $f_2^{(c)}$. Table 4.1 shows values² of this hypothetical asset for a range of single-index moves. Specifically, the second column was obtained by keeping the Euro STOXX 50 level f_2 unchanged and equal to $f_2^{(c)}$ while the S&P 500 index was shifted by the fractional amount shown under *Index Move*. Similarly, when calculating the third column, the S&P 500 level remained constant at $f_1 = f_1^{(c)}$ and f_2 changed. An inspection of this table reveals that the model valuation increases and decreases in line with changes in both the

¹Since vector f is a column, technically we should transpose the row: $f_{\text{current}} \equiv (f_1^{(c)}, f_2^{(c)})^T$, the superscript T signifying a transpose. In order to reduce clutter, we will omit the transpose symbol acting on vectors except when such an omission could result in ambiguity.

²While the discussions of this section are general and their conclusions apply to most assets, for the purposes of numerical illustration we used an imaginary financial instrument whose value equals $A[c_1^2 + c_2^2]^{1/2}$ where c_1 and c_2 are prices of European call options on, respectively, S&P 500 and the Euro STOXX 50, struck at the money and expiring in one year, and A is an inconsequential multiplier that forces the asset's value to equal \$100 million at spot index levels. Although admittedly contrived, this instrument saliently exhibits the characteristic properties of a complex asset-backed security, most notably the nonlinearity of factor dependence and the related non-additivity of responses to market shocks.

TABLE 4.1 Response Matrix for a Hypothetical Asset with Two Market Factors. Single-Index Moves

Index Move	S&P 500 (\$mn)	Euro STOXX 50 (\$mn)
−50%	69.86	71.55
−30%	70.05	71.76
−10%	80.11	81.41
−5%	88.46	89.22
−2%	95.01	95.34
−1%	97.45	97.61
nil	100.00	100.00
+1%	102.67	102.50
+2%	105.45	105.10
+5%	114.40	113.49
+10%	131.06	129.13
+30%	207.62	202.59
+50%	281.19	274.82

S&P 500 and the Euro STOXX 50, leading one to conclude that the asset has a positive exposure to both markets.

A response matrix like the one in Table 4.1 can furnish useful guidance to an asset's exposures and point out market moves that would endanger a portfolio. It can also direct a manager toward steps that can pre-empt adverse events—in this example, such steps could consist in hedging against a decline in U.S. and euro-zone stocks. As a quantitative guide, however, it is of little use except for small index changes, not more than a few percent. The reason is that, first, world markets are interconnected, and, for instance, a large drop in U.S. equities is almost certain to be accompanied by a commensurate decline in European stocks, and, second, assets typically depend on market factors in a nonlinear fashion and, consequently, the changes in asset value, resulting from moves in different factors, are *non-additive*. While it is plausible for the S&P 500 to rise 1% with the Euro STOXX 50 remaining unchanged or even declining a few percentage points during the same period, a 50% collapse in the S&P 500 with European stocks remaining unchanged is, to use an understatement, extremely unlikely. A manager can use Table 4.1 to estimate the change in portfolio value resulting from a 1% increase in the S&P 500 and a concurrent 2% drop in the Euro STOXX 50 (\$2.67 mn − \$4.66 mn = \$ − 1.99 mn), but it would be fallacious to infer that a 50% crash in the Euro STOXX 50 would result in a \$28.45 mn (= \$100 mn − \$71.55 mn) decline in portfolio value: This calculation ignores the almost certain concomitant decline in the S&P 500. Furthermore, because of the asset's nonlinearity, it would generally be wrong, having assumed that both

markets move by about the same fractional amount, to combine the losses implied by Table 4.1 for 50% drops in both indices (i.e., to estimate the asset loss to be \$30.14 mn + \$28.45 mn = \$58.59 mn).

Scenario Grid

The problems of correlated market moves and nonlinear factor dependence can be addressed by tabulating valuations for *simultaneous* market factor moves, that is, for a range of values of vector f as a whole. Table 4.2 shows valuations of our hypothetical asset for f equal to $(0.5 f_1^{(c)}, 0.5 f_2^{(c)})$, $(0.5 f_1^{(c)}, 0.7 f_2^{(c)})$, $(0.5 f_1^{(c)}, 0.98 f_2^{(c)})$, and so on. The entries in this *scenario grid* contain actual model valuations and offer quantitatively meaningful measures of the asset’s market risks. These valuations, however, are still available only for a discrete set of values of f . An attempt to refine the grid by adding more factor level points than the seven displayed in Table 4.2 is likely soon to run into computer resource limitations, especially in a realistic case of multiple market factors: If one chooses ν distinct levels for each of the K market factors, then the number of points on the K -dimensional grid equals ν^K , liable to be prohibitively large in realistic situations. (The scenario grid becomes a multidimensional cube for $K \geq 3$ —or “hypercube” for $K \geq 4$ —and can in practice be displayed as a collection of two-dimensional “slices.”) For example, for the modest values of $\nu = 7$ and $K = 4$, the generalization of the grid in Table 4.2 would require an unrealistic 2,401 asset valuations. Note that our sample hypothetical asset has in fact four market factors if we include the two relevant risk-free rates, one for the United States and another for the euro zone—and six if we count the indices’ volatilities.

The problem of computational intensity can be somewhat mitigated by computing a sparser grid and interpolating values at intermediate points. This procedure, however, can introduce complications of its own as

TABLE 4.2 Two-Dimensional Scenario Grid for a Hypothetical Asset with Two Market Factors. The unlikely scenarios are italicized. Asset values in \$mn

Euro STOXX 50 \ S&P 500	-50%	-30%	-2%	nil	+2%	+30%	+50%
-50%	0.1	5.2	64.4	71.6	79.0	195.5	272.4
-30%	5.5	7.6	64.6	71.8	79.2	195.6	272.4
-2%	63.0	63.2	90.1	95.3	101.0	205.4	279.6
nil	69.9	70.1	95.0	100.0	105.5	207.6	281.2
-2%	77.0	77.2	100.4	105.1	110.3	210.1	283.0
+30%	189.5	189.6	200.2	202.6	205.3	272.3	331.8
+50%	265.3	265.4	273.1	274.8	276.9	329.6	380.3

TABLE 4.3 Response Matrix with Concurrent Equally Sized Market Moves for a Hypothetical Asset. Bottom line shows the (incorrect) valuation achieved by combining the effects of single market moves from Table 4.1

Equity	-50%	-30%	-10%	-2%	nil	+2%	+10%	+30%	+50%
Model Value (\$mn)	0.1	7.6	55.2	90.1	100.0	110.3	154.4	272.3	380.3
Single-Market (\$mn)	41.4	41.8	61.5	90.4	100.0	110.5	160.2	310.2	456.0

multidimensional interpolation can be tricky (Press et al., 2002) and its accuracy unsatisfactory, unless the interpolated grid is sufficiently dense.

Our final comment on the scenario grid is that mindless computation of every grid point would be wasteful since many of the scenarios are rather unlikely. For instance, the italicized values in Table 4.2 correspond to highly improbable combinations of either large opposite moves in European and U.S. stocks or a large move in one of the two markets with the other practically unchanged. This observation leads to the issue of judicious *scenario selection*, on which we will say more in the rest of this subsection, and the related topic of *stress tests*, discussed in greater detail in section 4.3.

Synchronized Response Matrix

A computationally economical and practically useful way of assessing an asset's vulnerability to large market shocks is to create a matrix of responses to *synchronized* factor moves. In the simplest case, one can assume that all related equity markets, as an example, change by the same fractional amount. Table 4.3 presents a synchronized response matrix for our hypothetical asset. It also shows for comparison the results of a naïve arithmetic addition of the effects of single-market moves from Table 4.1. As expected, the addition yields tolerable agreement with the "exact" model valuation for small index moves of a few percent; for catastrophic market declines, the single-market estimate is completely, and dangerously, wrong.

Market volatility can have a profound effect on asset valuation. By ignoring the volatility changes accompanying index moves, we have greatly underestimated our hypothetical sample asset's value in declining markets. In Table 4.4, in addition to synchronizing the index moves, we also adjusted the volatility of index returns (see section 4.3 for the methodology of volatility adjusting). A comparison of the results with Table 4.3 reveals the dramatic impact of volatility on asset valuation, especially for large adverse market moves.

TABLE 4.4 Response Matrix for a Hypothetical Asset. Synchronized index (percent change from current) and volatility (volatility points shift from current) moves

Equity	−50%	−30%	−10%	−2%	nil	+2%	+10%	+30%	+50%
Vol	+25%	+15%	+5%	+1%	nil	−1%	−5%	−12%	−17%
Value (\$mn)	22.6	44.2	75.6	94.5	100.0	106.2	136.1	254.3	375.5

4.3 STRESS-TESTING

The results of Table 4.4, although incorporating the effect of market moves on volatility, still ignore the probable co-movement of other market and non-market factors. For instance, the value of our hypothetical asset would be affected by interest rates, which are likely to change in conjunction with large market moves. Depending on the nature of a security, its value can also depend on the unemployment level, mortgage prepayment speed, delinquency rates, credit ratings, market sentiment indices, and other factors. Creating a response matrix that reflects all the relevant parameters is in general impractical, and a risk manager would benefit far more from considering the impact on the asset of a selection of extreme events in which all the relevant factors co-move in a realistically coordinated way, that is, in a manner that reflects their historical co-movements and causal and co-variant relationships (of course, the risk manager should always be cognizant of the possibility of changes in these relationships). Constructing the special kind of response matrix, resulting from this analysis, is known as *stress-testing*. In contrast to conventional response matrices or scenario grids, stress tests do not strive to present a broad view of an asset's valuations for all conceivable eventualities but instead focus on a few carefully chosen scenarios, typically detrimental to the asset.

An analyst can decide on the magnitude of the relative co-movements of market factors on the basis of either qualitative or quantitative considerations (or a combination of the two). An example of the former would be a decision to consider the leftmost two diagonal scenarios in Table 4.2: if the S_{px} drops by 30 or even 50 percent, then the [estoxx] is also likely to go down by a similar, if not exactly the same, fraction. An analyst might be able to get a sense of the impact of a catastrophic market decline on the portfolio by considering simultaneous drops in the two indices of 10, 30, and 50 percent, that is the first three shocks in Table 4.4.

In some cases (e.g., an equity index and its volatility, or the prevalent interest rates) the relative magnitude of coincident moves can be not obvious. In this case an analyst can utilize *linear regression*, described in Appendix B,

to calculate the market factors' *betas*, which we will introduce in Section 5.2 (equations (5.11) and (5.12)), and then calculate the relative factor moves similarly to the procedure for determining an optimal hedge size, which we will describe in Section 5.3. As an example, in order to decide on volatility shifts in Table 4.4 we regressed historical changes in the indices' implied volatilities on the indices' returns.

4.4 SENSITIVITIES

Sensitivities are measures of an asset value's response to small changes in continuous factors. As in the case of response matrices, our main focus will be on market factors $f \equiv (f_1, \dots, f_K)$, although the discussion can be straightforwardly extended to continuous non-market ones.

Traditionally, most sensitivities are denoted with Greek letters and are collectively known as "the Greeks." A notable exception is sensitivity to volatility, known as the *vega*. Of course, Vega is a star in the constellation of Lyra rather than a letter of the Greek alphabet. It was also a model of motor cars, produced by Chevrolet in the 1970s, and rumored to have inspired Myron Scholes, an owner. Notwithstanding the lack of proper credentials, the vega is considered a member of the Greeks in the context of market exposures measurement; we will denote it with the symbol $\mathcal{V}$. Occasionally, conformity prevails and one encounters alternative names for sensitivity to volatility, which are *bona fide* Greek characters: lambda (λ), kappa (κ), and tau (τ). (To add confusion, lambda is increasingly used nowadays to refer to elasticity, defined in equation (4.14), namely, the fractional change in the price of a derivative for a fractional change in the underlying.)

Linear Sensitivities

Consider first the case of a single market factor whose current value is f , the corresponding asset value being $V(f)$. We want to know the change ΔV in the value of the asset if f changes by a small amount Δf (we assume $\Delta f \neq 0$ and will quantify the meaning of "small" later in this subsection). With the help of the asset model we can easily calculate $\Delta V = V(f + \Delta f) - V(f)$. Now let us rewrite this expression as

$$\Delta V = \frac{V(f + \Delta f) - V(f)}{\Delta f} \Delta f. \quad (4.2)$$

The magnitude of both the numerator and the denominator of the quotient in this equation becomes smaller as Δf (which can be negative, as well as positive) shrinks. However, the value of the quotient itself approaches a

constant, whose magnitude is not necessarily small and is *independent of* Δf . This statement is true for all functions $V(f)$ that can conceivably arise in quantitative finance for most values of the argument f ; it may not hold on a countable (and typically small—usually one) number of points—for instance, when the underlying price for an option at expiration equals the strike. Barring such isolated singular price points (usually for trivial and uninteresting cases of nonlinear derivatives at expiration), our assertion implies that the equality in equation (4.2) holds approximately even if we reduce the argument shift in the quotient:

$$\Delta V(f, \Delta f) \approx \frac{V(f + \epsilon \Delta f) - V(f)}{\epsilon \Delta f} \Delta f \quad (4.3)$$

where ϵ is an arbitrary positive number smaller than unity, $0 < \epsilon < 1$. We stress that this estimate of the change in V is for the *original* factor shift Δf , as emphasized by the explicitly stated argument of ΔV .

A practically significant consequence of our observation is that we can calculate the quotient in equation (4.3) just once, using any sufficiently small change in f , and then use the result to evaluate $\Delta V(f, \Delta f)$ for various factor shifts Δf . In fact, it is convenient to calculate the quotient in the limit of very small but non-zero (*infinitesimally* small) ϵ . We call the result the asset's *linear sensitivity* χ to factor f and, having determined χ , can evaluate ΔV for a range of factor shifts:

$$\Delta V(f, \Delta f) \approx \chi(f) \Delta f. \quad (4.4)$$

The sensitivity χ is mathematically a derivative of V with respect to f and is formally defined as

$$\chi(f) = \lim_{u \rightarrow 0} \frac{V(f + u) - V(f)}{u} \equiv \frac{dV}{df}. \quad (4.5)$$

The symbol “ $\lim_{u \rightarrow 0}$ ” means that one should evaluate the fraction for a very small but non-vanishing u ; for example, if one did so for a sequence of $u = f/10, f/100, f/1,000$, and so on, then, assuming $f \neq 0$, the results would converge to a number that we call χ .

For an asset, dependent on multiple market factors $f = (f_1, \dots, f_K)$, we can define the sensitivity to each of them analogously to equation (4.5):

$$\chi_k(f) = \lim_{u \rightarrow 0} \frac{V(f_1, \dots, f_k + u, \dots, f_K) - V(f)}{u} \equiv \frac{\partial V}{\partial f_k}, \quad k = 1, 2, \dots, K \quad (4.6)$$

where the partial derivative symbol ∂ signifies that only one, the k th, factor is being changed while the rest are kept constant. Generalizing equation (4.4), one can now estimate the change in V resulting from a shift in f as

$$\Delta V(f, \Delta f) \approx \sum_{k=1}^K \chi_k(f) \Delta f_k. \quad (4.7)$$

Readers with a quantitative background will recognize this approximation as the linear part of a Taylor expansion. In fact, the *Taylor theorem* of differential calculus (see, e.g., Courant, 1952) forms the mathematical underpinning for the use of factor sensitivities to measure asset value changes in response to continuous factor moves; in particular, it provides justification for adding the effects of individual factor shifts in equation (4.7) and can furnish an estimate of the approximation's error.

We reiterate that by their definition sensitivities are functions of f but *not* of the factor shift Δf . This property distinguishes sensitivities from valuation changes $\Delta V(f, \Delta f)$, which do depend on both the factor levels and the size of their shift.

The Greeks (Linear) In equation (4.6) we used a generic notation χ_k for all linear sensitivities to market factors. Historically, sensitivities to certain types of factors have been denoted with Greek letters. In time, the sensitivities themselves absorbed these letters' names and came to be known by them. We will now describe the most important sensitivities and introduce their conventional names.

Price Sensitivity: The Delta Perhaps the most commonly encountered, and probably the most important for the purpose of risk analysis and hedging, is the *delta*, the sensitivity to an underlying instrument's price p_i :

$$\delta_i = \frac{\partial V}{\partial p_i}. \quad (4.8)$$

Here p_i could be the price of a stock, a bond, an index, a commodity, a futures contract, or of any other investment instrument, determining the value V of the asset. In fact, it need not even be an actual price: In the case of a variance option, as an example, p_i is the variance of the reference quantity (e.g., price return on a stock), while for an interest rate derivative p_i is the rate. Complex asset-backed securities, dependent on multiple markets or

share prices, have multiple deltas, whereas an option on a stock, variance, or an interest rate has only one.

Deltas enable one to estimate the effect on the asset value of small price moves, all else staying constant:

$$\Delta V \approx \sum_i \delta_i \Delta p_i \quad (4.9)$$

(over prices)

where the summation extends over all the prices among the asset's market factors. In other words, the *absolute deltas*, defined in equation (4.8), can be viewed as conversion factors, translating dollar price changes of underlying instruments to corresponding dollar changes in the asset's value.

Relative Deltas Because absolute deltas convert dollars to dollars, comparing them would not in general reveal the relative significance of the corresponding exposures and, therefore, highlight areas of elevated risk that require special attention. This is especially true in the case of factors with significantly different typical values. Consider, as an example, a security with exposures to both the NASDAQ 100 Index and the S&P 500 Telecommunication Services Index and let the corresponding deltas be both equal to 0.5 (this security could, e.g., be a portfolio, consisting of one option contract on each index). A decline of 10 points in either index would lead to an approximate \$5 reduction in the security's valuation (we assume the value of one index point to be \$1). It would be incorrect, however, to conclude that the security's risk exposures to the two indices are the same. Indeed, whereas for the NASDAQ 100, which ended the year 2009 at 1,860, a 10-point reduction would represent a tame 0.5% decline, for the S&P 500 Telecoms, closing the year at 115, the same point move would constitute a substantial 9% drop. The changes in the two indices are unlikely to be so disparate during the same time period, and if the decline in the NASDAQ 100 were a commensurate 9%, then the loss, resulting from the NASDAQ 100 exposure, would be a much more conspicuous $1,860 \times 9\% \times 0.5 = \84 : the security's risk exposure to the NASDAQ 100, is, in fact, much greater than to the S&P 500 Telecoms—absolute deltas by themselves are poor indicators of relative risk exposures.

A sensitivity measure, comparable across factors, would be one that shows an asset's reaction to a given *fractional* price change. Let us rewrite equation (4.9) as

$$\Delta V \approx \sum_i [p_i \delta_i] \times \frac{\Delta p_i}{p_i} \equiv \sum_i \hat{\delta}_i R_i. \quad (4.10)$$

(over prices) (over prices)

Here we introduced the *relative delta*

$$\hat{\delta}_i = p_i \delta_i \equiv p_i \frac{\partial V}{\partial p_i} \quad (4.11)$$

that converts fractional price change $R_i \equiv \Delta p_i / p_i$ to the dollar change in the asset value. Readers familiar with differential calculus may have observed that the relative delta is a logarithmic derivative:

$$\hat{\delta}_i = \frac{\partial V}{\partial \ln p_i} \approx \frac{\Delta V}{\Delta \ln p_i} \quad (4.12)$$

where the small finite difference $\Delta \ln p_i \equiv \ln[p_i + \Delta p_i] - \ln p_i = \ln[(p_i + \Delta p_i)/p_i]$ is, in fact, a logarithmic return, defined in footnote 16 of Chapter 3. Writing $\hat{\delta}_i$ in this form makes explicitly manifest the fact that it is a sensitivity to price *returns*,³ in contrast with the absolute delta, which, as is clear from equation (4.8), is a sensitivity to price *changes*.

Returning to our example, we can calculate the sample security's $\hat{\delta}_i$ s: $1,860 \times 0.5 = \$930$ for the NASDAQ 100 and $\$115 \times 0.5 = \58 for the S&P 500 Telecoms exposure. An examination of these *relative* deltas, unlike absolute ones, instantly tells the manager that the security's exposure to the NASDAQ 100 is significantly greater than to the S&P 500 Telecoms and, therefore, should probably receive more attention.

It is common to express relative deltas as the asset value's change in response to a 1% increase in price, that is, for $R_i = 0.01$. For instance, our sample asset's response to a 1% jump in the NASDAQ 100 would be a \$9.30 rise in value, and to an equally sized shift in the S&P 500 Telecoms, a \$0.58 rise.

³Note that, as mentioned in footnote 16 of Chapter 3, the distinction between logarithmic and arithmetic returns vanishes in the limit of small Δp_i , and specification of the type of return is, therefore, unnecessary. In fact, approximating the derivative in equation (4.11) with a ratio of small finite differences, one could directly express the relative delta as a sensitivity to (small) arithmetic returns: $\hat{\delta}_i \approx \Delta V / (\Delta p_i / p_i)$.

Equivalent Variation Investment In section 2.5 we introduced *equivalent variation investment* (EVI), a quantitative measure of market exposure, and demonstrated on an example how to use it to size a hedge. We define EVI as *the size of an investment in a market factor instrument (stock, bond, index, etc.) such that the variations in this investment's value with small factor moves equal the asset's value changes, attributable to these moves.*

To understand the meaning and significance of EVI, consider an asset, dependent on I prices p_i ($i = 1, \dots, I$; $I \leq K$, where K is the total number of market factors), and let the corresponding EVIs be $\mathcal{E}_i$. For each price p_i let us invest in its underlying factor instrument an amount $\mathcal{E}_i$. Then the changes in the value of the resulting portfolio $H = \{\mathcal{E}_1, \dots, \mathcal{E}_I\}$ would equal the changes in asset value for arbitrary (small) changes in prices p_i , other, non-price, factors remaining fixed. The corollary is that *selling short* portfolio H would compensate for changes in asset value, that is, would provide a *hedge* against price moves.

This hedge, of course, would be imperfect: The calculation of EVIs would be subject to conceptual and formalism model uncertainties, defined in section 3.2 (the consequent mismatch between the asset's calculated and actual behavior would contribute to basis risk). In addition, for larger price moves, the performance of portfolio H and of the asset would in general diverge because of the asset's nonlinear factor dependence.

We further discuss the issue of hedge size selection in relation to EVI in section 5.3 and in Chapter 6. Here we make an important observation that EVI equals the relative delta:

$$\mathcal{E}_i = \widehat{\delta}_i \equiv p_i \frac{\partial V}{\partial p_i}. \quad (4.13)$$

The equality follows from equation (4.10): If one invested an amount $\mathcal{E}_i = \widehat{\delta}_i$ in factor instrument i , and this instrument subsequently delivered a rate of return R_i , then the value of this investment would change by $\mathcal{E}_i R_i$. This is exactly the portion $\widehat{\delta}_i R_i$ of the asset value change ΔV_i , attributed, according to equation (4.10), to the move in price p_i . Simplifying somewhat and ignoring basis risk, one can say that as far as small price changes are concerned, holding the asset is equivalent to holding portfolio H .

Returning once more to our sample security, we claim that holding it would be approximately equivalent to investing \$930 in the NASDAQ 100 and \$58 in the S&P 500 Telecoms. The equivalence pertains for small index moves and under the assumption of no changes in, for instance, index volatilities and in interest rates. A manager could hedge his exposure to these two indices by shorting \$930 of the NASDAQ 100 and \$58 of the

S&P 500 Telecoms. (The manager should also consider hedging volatility, interest rate, and other non-price exposures with other instruments.)

Elasticity: The Omega While the relative delta translates a fractional price change to the asset value dollar change, *elasticity*, or *omega*, extends the concept by converting fractional price change to the *fractional* change in asset value:

$$\Omega_i = \frac{p_i}{V} \times \frac{\partial V}{\partial p_i} = \frac{\partial \ln V}{\partial \ln p_i}. \quad (4.14)$$

The second equality expresses elasticity as *return* on the asset generated by a unit return on price p_i —see discussion following equation (4.12).

Elasticity is also known as *lambda*, λ_i .

Volatility Sensitivity: The Vega Tables 4.3 and 4.4 illustrate that a change in volatility can profoundly impact asset value. The sensitivity to changes in *standard deviation* σ_i of (logarithmic) returns on price p_i is known as the *vega*:

$$\mathcal{V}_i = \frac{\partial V}{\partial \sigma_i}. \quad (4.15)$$

While being, as already observed, an impostor in the midst of Grecian alphabetic characters, the vega has firmly established itself as one of the “Greeks.” Vegas are usually expressed as assets’ dollar value change in response to a one-point increase in volatility (we assume the volatility to be annualized, as discussed in Appendix B).

Often, especially in the context of volatility derivatives, such as variance swaps and variance options the appropriate quantity is the sensitivity to the *variance* of returns σ_i^2 , the *variance vega*, or *upsilon*:

$$\Upsilon_i = \frac{\partial V}{\partial (\sigma_i^2)}. \quad (4.16)$$

The two measures of responsiveness to volatility changes are related by a simple identity:

$$\mathcal{V}_i = 2\sigma_i \Upsilon_i. \quad (4.17)$$

As already mentioned, the “price” p_i in equations (4.8) and (4.11) is not necessarily an actual price or index level. For example, the σ_i^2 in equation

(4.16) could be, in the case of a variance option or swap, a *variance of variance*. Similarly, the σ_i in equation (4.15) could stand for interest rate volatility if p_i were a (forward) interest rate.

Interest Rates and the Rho The value of many complex assets depends on multiple interest rates, which can enter the valuation process either as market factors or through the discount curve. To complicate matters, the distinction may become blurred when risk-neutral valuation is used. For example, the price of a plain vanilla stock or index option, valued within the Black-Scholes-Merton framework, is determined by a single rate r_f , which serves as both the required rate of return (a market factor) and the discount rate. The number of distinct rates involved multiplies with the increasing complexity of the asset. For instance, even for a simple foreign exchange option there are already two risk-free rates, one for the “foreign” and one for the “domestic” (or base) currency.

Interest rates (other than a single r_f , as in the option example) that enter a model as market factors are typically forward rates, serving as underlying instruments for interest rate derivatives. For instance, an interest rate option involves a single forward rate. Another example is a floating rate bond, which one can view as a derivative with multiple underlying forward rates, one for each coupon date.

Somewhat less frequently, models can involve market-determined expected rates of return. Rate expectations can crop up as a result of a discontinuity assumption (e.g., the jump-diffusion process discussed at the end of section 3.3) or of dispensing with the risk-neutral valuation principle. There is no reason to treat these rates differently than other market factors, and an asset value’s sensitivity to them is usually termed “delta” (to a degree this is a matter of semantics).

When speaking explicitly of “interest rate sensitivities,” we will mean sensitivities with respect to discount rates. As mentioned, complex assets as a rule depend on multiple rates. This multiplicity can derive from exposure to several currency markets, as in the foreign exchange option used in the example above, or may reflect different maturities, exemplified by the above-mentioned floating rate bond. Of course, a single asset can combine both sources of multiplicity.

Spurred by the especial importance of interest rates for debt instruments and by the complexity of debt instruments’ dependence on the rate term structure, fixed income market participants have devised an array of quantitative measures that probe various aspects of interest rate risk: duration PV01, convexity, key rate measures; we discuss these “fixed-income-style” characteristics (usually not considered members of the Greeks) under a separate heading below.

In the context of derivatives valuation the sensitivity to the interest rate is known as the *rho*. For a security, dependent on a single constant risk-free rate r_f (e.g., a plain vanilla option), the rho is defined as

$$\rho = \frac{\partial V}{\partial r_f} \quad (4.18)$$

(as already noted after equation (4.6), the partial derivative, indicated by the symbol “ ∂ ,” means that all parameters other than r_f should be kept fixed). It is easy to generalize this definition to an asset, whose valuation involves a discount curve $r_f(t)$, t being time to maturity, by replacing a change in the single risk-free rate with a *parallel shift* of the whole curve:

$$\rho = \left. \frac{\partial V}{\partial \Delta r_{\parallel}} \right|_{\Delta r_{\parallel}=0} \quad (4.19)$$

where the shifted curve is given by $r_f(t) + \Delta r_{\parallel}$ and the derivative is evaluated at $\Delta r_{\parallel} = 0$.⁴ For an asset with multiple currency exposures, a separate rho for each currency could be defined similarly.

Time Decay: The Theta The *theta*, or *time decay*, quantifies the asset value loss resulting from the passage of time:

$$\theta = \frac{\partial V}{\partial t}, \quad (4.20)$$

conventionally expressed as value change per calendar day. Most derivatives’ and asset-backed securities’ value decays with time if factors and rates do not change and, therefore, the theta is usually negative. There are, however, exceptions. For example, an in-the-money European put option on a stock can have a positive theta, that is, this option would grow in value with time. This makes intuitive sense: The shorter the time to expiration, the less likely the option is to expire out of the money. At the same time, the potential for appreciation is limited: The price of the underlying stock cannot become negative. A European call option’s theta, by contrast, would always

⁴This means that, having calculated the derivative as a function of $\Delta r_{\parallel}$ as

$$\lim_{u \rightarrow 0} \frac{V(\Delta r_{\parallel} + u) - V(\Delta r_{\parallel})}{u}$$

a special case of equation (4.5), one sets $\Delta r_{\parallel} = 0$ in the result.

be negative because, no matter how deep in the money the option is, there is a potential, diminishing with time, for further unlimited appreciation.

Beware that sometimes the theta is defined with an opposite sign, $-\partial V/\partial t$; for most assets this quantity is positive and equals the magnitude of daily loss to time decay.

Unlike delta, vega, or rho, one cannot hedge theta. The nature of hedging is to reduce the effect of uncertainty, be it of future evolution of price, volatility, or interest rates. There is no uncertainty about the passage of time. It is, nonetheless, useful for an asset manager to know how much his portfolio will lose in a given period of time if markets remain stagnant. A frequently used related quantity is the *theta breakeven move*, the amount by which the underlying market factor or factors need to change, usually in one day, to offset the effect of time decay. While there exists no theta hedge in the literal sense, the manager often can adjust a portfolio's composition to reduce time decay, usually at the expense of some desired characteristic (e.g., gamma).

Note that although the rho and the theta are not sensitivities to market (nor non-market, for that matter) factors, we still consider them special cases of the generic sensitivity χ_k , *pace* the punctilious reader.

Second Order Sensitivities: Gammas

Equations (4.7) and (4.10) use linear sensitivities to estimate asset value change in response to market factor moves. This estimate can be accurate for small shifts but, owing to its failure to capture the usually present non-linearity of the asset response, it deviates from model valuation as the size of factor moves increases. The deterioration of the approximation accuracy is illustrated in Figure 4.1, which compares the model valuation of a call option with a linear approximation applied at the current spot price: Because of the curvature of the option price's dependence on the stock price, the error becomes too large for the estimate to be useful when the underlying price changes by more than a few percent.

What does "linearity" mean? It implies that the value is a linear function of market factors, namely, focusing for simplicity on a single factor f , that $V(f) = Af + B$ where A and B are some *constants* (i.e., numbers, independent of f). In the linear case, the sensitivity $\chi \equiv dV/df = A$ is itself a constant, independent of f . If the option were a linear instrument, then the exact model price would have been the straight (dashed) line in Figure 4.1. The deviation arises because the actual sensitivity (in our example the option delta) varies with f , the underlying stock price. The delta is in fact the slope of the price curve (i.e., of a tangent straight line) at a given point, and the fact that the model curve's slope varies signifies nonlinearity.




FIGURE 4.1 Comparison of model call option price with linear and quadratic approximations. Results shown are for an at-the-money European call on a stock when the underlying spot price is \$100.

This insight suggests a way to reduce the approximation error by adjusting for the fact that the sensitivity is not constant over the interval Δf . One way of doing so is to replace $\chi(f)$ in equation (4.4) with the average of its values at the end points of the interval:

$$\begin{aligned}
 \Delta V(f, \Delta f) &\approx \frac{1}{2}[\chi(f) + \chi(f + \Delta f)]\Delta f \\
 &= \frac{1}{2}[2\chi(f) + (\chi(f + \Delta f) - \chi(f))]\Delta f \\
 &= \chi(f)\Delta f + \frac{1}{2} \frac{\chi(f + \Delta f) - \chi(f)}{\Delta f} \Delta f^2. \quad (4.21)
 \end{aligned}$$

Now we can apply to the quotient in the second term the same argument as when deriving equation (4.4) from equation (4.2): As we shrink the size of the shift Δf *in the quotient only*, the quotient tends to a value, independent of Δf . This limiting value is the “sensitivity of sensitivity,” or, mathematically,

the derivative of χ with respect to f or the *second derivative* of V , and is known as the *gamma*:⁵

$$\Gamma(f) = \frac{d\chi}{df} = \frac{d^2V}{df^2}. \quad (4.22)$$

Knowing the gamma, one can refine the estimate in equation (4.4):

$$\Delta V(f, \Delta f) \approx \chi(f) \Delta f + \frac{1}{2} \Gamma(f) \Delta f^2. \quad (4.23)$$

The resulting dramatic improvement in the approximation accuracy is evident in Figure 4.1: The quadratic (dash-dotted line) approximation stays meaningfully close to the model price curve for price changes of up to tens percent.

For an asset with multiple factors, equations (4.22) and (4.23) generalize as follows:

$$\Gamma_{kk'}(f) = \frac{\partial \chi_k}{\partial f_{k'}} = \frac{\partial^2 V}{\partial f_k \partial f_{k'}}, \quad (4.24)$$

$$\Delta V(f, \Delta f) \approx \sum_k \chi_k(f) \Delta f_k + \frac{1}{2} \sum_{kk'} \Gamma_{kk'}(f) \Delta f_k \Delta f_{k'} \quad (4.25)$$

where summation over all the market factors is understood, that is, over $k = 1, \dots, K$ and $k' = 1, \dots, K$. Mathematically, this expression extends the Taylor expansion in equation (4.7) to quadratic terms.

The inclusion of quadratic terms engenders, in the case of multiple factors, a qualitatively new type of object, the *cross* sensitivity, that is, the sensitivity $\Gamma_{kk'}$ of χ_k to the change in a factor $f_{k'}$ with $k' \neq k$. Using the example defined in footnote 2 of a security dependent on the S&P 500 and the Euro STOXX 50, the implication of a non-vanishing cross gamma is that a move in say the Euro STOXX 50 would change the delta δ_1 with respect to the S&P 500, even if the U.S. equity index did not move.

Note that with the possible exception of a few special price points for nonlinear securities, typically at expiration, second-order sensitivities are *symmetric*, that is, the order of their indices does not matter: $\Gamma_{kk'} = \Gamma_{k'k}$.

⁵The term “gamma” is traditionally applied to second-order sensitivities with respect to prices or price-like market factors. We use this term generically.

The change in deltas and other sensitivities is approximated by

$$\Delta \chi_k \approx \sum_{k'=1}^K \Gamma_{kk'} \Delta f_{k'}. \quad (4.26)$$

Diagonal price gammas $\Gamma_{ii} \equiv \partial \chi_i / \partial p_i = \partial^2 V / \partial p_i^2$ are the ones most commonly quoted, and they possess some special features. For a given i , the diagonal gamma measures the rate of change of sensitivity χ_i to p_i with respect to the change in the same price p_i . The sign of a diagonal gamma is apparent from the graph of the asset's value versus the corresponding factor: at points where the graph is downwards convex, the gamma⁶ is positive; at points of upward curve convexity, it is negative. For example, the model price curve in Figure 4.1 is everywhere downwards convex and, therefore, the gamma is everywhere positive (as it is for all plain vanilla options). Figure 4.1 also illustrates the benefits of holding positive gamma: The option price's response to stock price change is asymmetric in the owner's favor. Thus, a stock price increase of \$25 induces an option price rise of \$18, while an equal stock price drop causes the option to lose only \$9—positive diagonal gamma is generally good. Conversely, a negative diagonal gamma is invariably bad. Not surprisingly, assets with higher gamma are generally more expensive. As noted in section 2.5, in the context of delta hedging, option premium is in essence the cost of gamma.

Substantial gamma, even positive, can, nevertheless, be undesirable in a delta-hedged portfolio: Since gamma is the delta's sensitivity to market changes, a delta-neutral portfolio with a high gamma will fast lose its delta-neutrality as markets move, thus requiring frequent re-hedging and incurring transaction costs. To avoid this managers sometimes *gamma-hedge* their portfolios, in addition to delta-hedging, that is, they deliberately keep the magnitude of the portfolio's gamma low.

The Smallness Criterion The concept of gamma enables one to quantify the notion of “small” factor changes we have used throughout, especially when discussing linear approximations to value changes. A linear approximation begins to break down when linear terms become comparable to quadratic terms, that is when, for a given price p_i ,

$$|\delta_i \Delta p_i| \approx \frac{1}{2} |\Gamma_{ii}| \Delta p_i^2 \quad (4.27)$$

⁶When the meaning is clear from the context, we will often use “gamma” to mean “diagonal gamma.”

(vertical bars denote absolute value). This leads to the *smallness criterion*:

$$|\Delta p_i| \lesssim \left| \frac{\delta_i}{\Gamma_{ii}} \right|. \quad (4.28)$$

We stated the criterion for the most commonly encountered case of price sensitivities (deltas) but it applies for any continuous factor.

Second Order Greeks Nomenclature

As noted in footnote 5, the term “gamma” and the symbol Γ are usually reserved for price sensitivities of deltas. We, nevertheless, have used gamma generically to mean any second-order sensitivity. Here we list some commonly used names for specific second-order sensitivity measures. Like the vega, although considered members of the Greeks, most of these names are not Greek letters, and some not even words in any language.

Gamma (Narrow Sense) To reiterate, in the narrow sense *gammas* are price sensitivities of deltas:

$$\Gamma_{ii'} = \frac{\partial \delta_i}{\partial p_{i'}} = \frac{\partial \delta_{i'}}{\partial p_i} = \frac{\partial^2 V}{\partial p_i \partial p_{i'}}. \quad (4.29)$$

The order of the indices (e.g., whether one considers variations in the sensitivity to price p_i resulting from a change in price $p_{i'}$ or *vice versa*) is almost always unimportant. It may matter only at the points where the price surface $V(p_1, \dots, p_I)$ exhibits a kink, similar to an option price’s behavior at expiration when the underlying price equals the strike. This comment applies to all higher-order sensitivities.

Charm Also known as *delta decay*, *charm* is a cross-sensitivity between price and time. It can be viewed as either a delta’s sensitivity to time or the theta’s sensitivity to price:

$$\text{Charm}_i = \frac{\partial \delta_i}{\partial t} = \frac{\partial \theta}{\partial p_i} = \frac{\partial^2 V}{\partial t \partial p_i}. \quad (4.30)$$

DvegaDtime This quantity measures the vega's change with the passage of time; it can also be viewed the theta's sensitivity to volatility:

$$\text{DvegaDtime}_i = \frac{\partial \mathcal{V}_i}{\partial t} = \frac{\partial \theta}{\partial \sigma_i} = \frac{\partial^2 V}{\partial t \partial \sigma_i}. \quad (4.31)$$

Vanna A vega's sensitivity to price or, alternatively, a delta's sensitivity to volatility, is known as the *vanna*:

$$\text{Vanna}_{ii'} = \frac{\partial \mathcal{V}_i}{\partial p_{i'}} = \frac{\partial \delta_{i'}}{\partial \sigma_i} = \frac{\partial^2 V}{\partial \sigma_i \partial p_{i'}}. \quad (4.32)$$

Analysts are usually interested in the *diagonal* vanna

$$\text{Vanna}_i = \frac{\partial \mathcal{V}_i}{\partial p_i} = \frac{\partial \delta_i}{\partial \sigma_i} = \frac{\partial^2 V}{\partial \sigma_i \partial p_i}. \quad (4.33)$$

Volga A vega's sensitivity to volatility, known as *volga*, *vega convexity*, or *volatility gamma* (whence the contraction “volga” derives), is the value's second derivative with respect to the standard deviation of price returns:

$$\text{Volga}_{ii'} = \frac{\partial \mathcal{V}_i}{\partial \sigma_{i'}} = \frac{\partial^2 V}{\partial \sigma_i \partial \sigma_{i'}}. \quad (4.34)$$

As in the case of vanna, the diagonal version is the one most commonly used:

$$\text{Volga}_i = \frac{\partial \mathcal{V}_i}{\partial \sigma_i} = \frac{\partial^2 V}{\partial \sigma_i^2}. \quad (4.35)$$

The volga is also known as *vomma*, although the authors find the Russian river's name aesthetically preferable.

One should not confuse the volga with upsilon, or variance vega, equation (4.16), the asset value's first order sensitivity to *variance*.

Higher-Order Sensitivities

Given the remarkable improvement, evidenced in Figure 4.1, in the approximation accuracy achieved by incorporating gamma, a natural inclination might be not to stop at second-order sensitivities but to proceed to the third

and perhaps higher orders. The limit to this progression is set by the increasing model uncertainties, resulting from error amplification in sensitivity calculation, as we will explain shortly. These uncertainties can quickly overwhelm the estimated higher sensitivity values and any nominal improvement in valuation approximation. In addition, incorporation of higher orders in a Taylor expansion raises a mathematically nontrivial issue of the domain of convergence in Δf : Outside this domain, adding more terms in the Taylor expansion may actually *increase* the approximation error, sometimes dramatically.

Third-order sensitivities are sometimes monitored in the context of derivatives trading, especially of options. We list below the names and definitions of the most common measures, even though they are rarely used for complex asset-backed securities.

Color This is the rate of time decay of gamma:

$$\text{Color}_{ii'} = \frac{\partial \Gamma_{ii'}}{\partial t} = \frac{\partial^3 V}{\partial t \partial p_i \partial p_{i'}}. \quad (4.36)$$

Of greatest interest, and most commonly used, is the *diagonal* color:

$$\text{Color}_i = \frac{\partial \Gamma_{ii}}{\partial t} = \frac{\partial^3 V}{\partial t \partial^2 p_i}. \quad (4.37)$$

Given the stated importance of gamma in portfolio management and hedging, managers benefit from knowing the rate at which their assets lose or gain gamma as time progresses.

Speed While color measures the rate of gamma's time decay, *speed* is gamma's rate of change with price moves:

$$\text{Speed}_{ii''} = \frac{\partial \Gamma_{ii''}}{\partial p_{i''}} = \frac{\partial^3 V}{\partial p_i \partial p_{i'} \partial p_{i''}}. \quad (4.38)$$

As with color, the most commonly used (and, usually, the most useful) is the diagonal speed, namely, the sensitivity of a diagonal gamma to the corresponding price change:

$$\text{Speed}_i = \frac{\partial \Gamma_{ii}}{\partial p_i} = \frac{\partial^3 V}{\partial p_i^3}. \quad (4.39)$$

Ultima Also known as *DvolgaDvol*, this sensitivity measures the change in volga in response to a shift in volatility:

$$\text{Ultima}_{ii'i''} = \frac{\partial \text{Volga}_{ii'}}{\partial \sigma_{i''}} = \frac{\partial^3 V}{\partial \sigma_i \partial \sigma_{i'} \partial \sigma_{i''}}. \quad (4.40)$$

As in the case of color or speed, the most commonly used is diagonal ultima, that is, the one where all three of the involved volatilities characterize the same factor:

$$\text{Ultima}_i = \frac{\partial \text{Volga}_{ii}}{\partial \sigma_i} = \frac{\partial^3 V}{\partial \sigma_i^3}. \quad (4.41)$$

Zomma This last on our list of the Greeks, sometimes called *Dgamma-Dvol*, measures the gamma's rate of change with volatility, or the vanna's sensitivity to price:

$$\text{Zomma}_{ii'i''} = \frac{\partial \Gamma_{ii'}}{\partial \sigma_{i''}} = \frac{\partial \text{Vanna}_{ii''}}{\partial p_{i'}} = \frac{\partial^3 V}{\partial p_i \partial p_{i'} \partial \sigma_{i''}} \quad (4.42)$$

or, in the diagonal case,

$$\text{Zomma}_i = \frac{\partial \Gamma_{ii}}{\partial \sigma_i} = \frac{\partial \text{Vanna}_{ii}}{\partial p_i} = \frac{\partial^3 V}{\partial p_i^2 \partial \sigma_i}. \quad (4.43)$$

4.5 INTEREST RATE SENSITIVITIES: DURATION, PV01, CONVEXITY, KEY RATE MEASURES

In the preceding sections we introduced the general principles behind factor sensitivities and their uses. We also listed and discussed some commonly used factor sensitivity measures. When discussing interest rates, we made a distinction between rates that enter as market factors, sensitivities that are conceptually, and semantically, not different from other deltas, and those—rhos—that affect asset valuation via the discount curve. Here, as in the discussion of the rhos, we will focus on sensitivities to discount rates.

For many standard equity derivatives, such as options, the effect of interest rate shifts is usually less important than that of other factors, most notably the price and volatility of the underlying instrument. For these derivatives a single sensitivity to a parallel shift in yield curve usually suffices—this is the rho defined in equations (4.18) and (4.19). Interest rates typically play

a much more important role in valuation of fixed income and asset-backed securities. Moreover, these types of instrument are often intricately sensitive to the interest rate *term structure* and, in addition, their dependence on rates is frequently markedly nonlinear. Market practitioners use a number of sensitivity measures peculiar to interest rates, which are designed to capture these specificities.

Although traditionally considered separately and not viewed as members of the “equity-style” Greeks, interest rate sensitivities conceptually belong to the same category of asset characteristics that quantify responses to small factor moves.

Duration

We introduced duration in section 2.2, equation (2.1). Here we will extend its definition and explore different kinds of duration.

Modified and Macaulay Duration Consider an asset valued using discounted cash flows, as discussed in section 3.3. Equation (3.9) expresses the asset’s value in terms of cash flows and discount rates with annual compounding. Let us rewrite this expression for arbitrary compounding and let us, without loss of generality, set $t_0 = 0$, that is, let t_i be time to the i th cash flow. In addition, we shift the discount curve $r_f(t)$ by an amount $\Delta r_{\parallel}$, as we did in equation (4.19):

$$V(\Delta r_{\parallel}) = \sum_{i=1}^N \frac{C(t_i)}{(1 + (r_f(t_i) + \Delta r_{\parallel})/n)^{nt_i}} \quad (4.44)$$

where the interest rates are compounded n times per annum (e.g., $n = 1$ for annual compounding, 2 for semiannual, 4 for quarterly, and 12 for monthly).

Modified duration is minus scaled sensitivity of the valuation to the parallel shift $\Delta r_{\parallel}$ *assuming cash flows remain unchanged*:

$$\mathcal{D}_{\text{mod}} = -\frac{1}{V} \left. \frac{\partial V}{\partial \Delta r_{\parallel}} \right|_{\Delta r_{\parallel}=0} = \frac{1}{V} \sum_{i=1}^N t_i \frac{C(t_i)}{(1 + r_f(t_i)/n)^{nt_i+1}}. \quad (4.45)$$

Recalling the definition in equation (4.19), modified duration equals minus scaled rho:

$$\mathcal{D}_{\text{mod}} = -\frac{1}{V} \rho. \quad (4.46)$$

The minus sign makes the duration positive: If cash flows are unaffected by changes in discount rates, as is assumed in calculating modified duration, then it follows from equation (4.44) that the valuation always decreases with rising rates. (We assume that we are dealing with a “real” security, all of whose cash flows are positive, as opposed to, say, project valuation, in which case some of the $C(t_i)$ can be negative cash outflows.) Readers who have had exposure to differential calculus can easily derive equation (4.45) from (4.44); we ask the rest to accept the result on faith.

The expression for duration in equation (4.45) is *almost* the weighted average of the lengths of time until receipt of cash flows with these cash flows’ present values used as weightings:

$$\mathcal{D}_{\text{mod}} = \sum_{i=1}^N \frac{t_i}{1 + r_f(t_i)/n} \times \frac{\text{PV}(C(t_i))}{V}, \quad V = \sum_{i=1}^N \text{PV}(C(t_i)). \quad (4.47)$$

The superfluous factors $1/(1 + r_f(t_i)/n)$ tend to unity, and therefore become insignificant, with increasing compounding frequency (i.e., as n becomes large) or in a low-rates environment: For large n or small discount rates $r_f(t_i)$, modified duration is approximately the weighted average time till bond payouts. In fact, an alternative, and older, approach, proposed by Frederick R. Macaulay in 1938, is to define duration *exactly* as a weighted average time to cash flows:

$$\mathcal{D}_{\text{Macaulay}} = \sum_{i=1}^N t_i \times \frac{\text{PV}(C(t_i))}{V}. \quad (4.48)$$

The meaning of this *Macaulay duration*⁷ can be clarified by considering a zero-coupon bond, that is a bond with a single cash flow $C(t_1)$ paid after a time period t_1 . In this case, the discounted cash flow (DCF) value is $V = \text{PV}(C(t_1))$, and the Macaulay duration equals exactly the time to maturity, that is, to the first and only cash flow: $\mathcal{D}_{\text{Macaulay}} = t_1$ (this argument, of course, ignores credit risk, i.e., the possibility of the bond’s defaulting prior to maturity). A coupon-paying bond, or any other DCF-valued asset, can

⁷In bond markets, when calculating duration, it is common to discount all the cash flows with the bond’s yield to maturity and to calculate duration with respect to changes in this yield—this is a more practical if conceptually less satisfactory approach. When the present values are calculated with a zero-coupon discount curve, as in equation (4.45), then the quantity in equation (4.48) is sometimes called the *Macaulay-Weil* or *true* duration.

be viewed as a collection of zero-coupons, and the Macaulay duration can be interpreted as these zero-coupon bonds' times to maturity weighted with the bonds' respective values. The zero-coupon analogy also demonstrates that, as long as cash flows are rate-independent, a coupon-bearing bond's Macaulay duration is shorter than its time to maturity.

The concept of duration was, as the name implies, originally introduced as a measure of the typical length of time over which an investor receives pay-outs from a bond—note in this respect that by definition duration is measured in units of time, typically years: The rates in equations (4.44) and (4.45) have units of inverse time (percent interest paid per annum), as does n (the number of compoundings in a year). Intuitively, the longer an investor is exposed to an asset, that is, the longer the asset's duration, the more likely the asset's value is to be affected by changes in prevailing interest rates. Consequently, duration is a measure of an asset's rate sensitivity. The fact that the duration of a coupon-paying bond is shorter than its time to maturity reflects the cash flows (coupon payments) that the investor will receive before the bond matures.

Macaulay duration, while having a clear intuitive interpretation, is not, however, an exact measure of interest rate sensitivity. That is why market participants *modified* the Macaulay duration to arrive at the exact sensitivity to parallel shifts in the yield curve, as defined in equations (4.45) and (4.47). For small parallel shifts, $\Delta r_{\parallel}$ modified duration allows one to evaluate the corresponding fractional change in asset value:

$$\frac{\Delta V}{V} \approx -\mathcal{D}_{\text{mod}} \Delta r_{\parallel} \quad (4.49)$$

with the approximation error vanishing in the limit of very small $\Delta r_{\parallel}$. In order to appreciate the difference between the two duration measures, note that while Macaulay duration can in general be used to obtain a meaningful approximation to the value change, $\Delta V/V \approx -\mathcal{D}_{\text{Macaulay}} \Delta r_{\parallel}$, especially in a low-rate environment, this approximation does *not* approach the exact relationship, no matter how small the shift $\Delta r_{\parallel}$ (as long as the number of compoundings n is finite, i.e., the definitions of modified and Macaulay durations are distinct—see equation (4.50)).

Macaulay duration has largely historical significance and is rarely used by modern practitioners. We will often use the term “duration” meaning “modified duration” or “effective duration,” which we define shortly (and which should have been consistently termed “effective modified duration”). Modified and Macaulay durations have identical information content and are related by a collection of trivial factors whose very presence derives from the artifice of finite compounding: They become unity in the limit

of continuous compounding (i.e., as $n \rightarrow \infty$) when the distinction between modified and Macaulay durations vanishes:

$$\mathcal{D}_{\text{mod}} = \mathcal{D}_{\text{Macaulay}} = \frac{1}{V} \sum_{i=1}^N t_i C(t_i) e^{-r_f(t_i)t_i} \quad (\text{continuous compounding}). \quad (4.50)$$

When present values are calculated with the bond's yield to maturity y or, alternatively, when the discount curve is flat (i.e., $r_f(t_i) = y$ for all i), then the relationship between the two durations reduces to a single factor:

$$\mathcal{D}_{\text{Macaulay}} = (1 + y/n) \mathcal{D}_{\text{mod}} \quad (\text{flat discount curve}). \quad (4.51)$$

Effective Duration When defining duration in equation (4.45), we assumed that the asset's sensitivity to interest rates enters only through the discounting process, that is, that cash flows $C(t_i)$ were unaffected by interest rate changes. This may be a reasonable assumption for government or high-grade corporate bullet bonds but not for bonds with embedded optionality (e.g., callable, puttable, or convertible), nor for many asset-backed securities.

A bullet bond's cash flows can change only as a consequence of default. Such an outcome is very unlikely (although not impossible, as discussed in section 2.2) for debt issued by developed economies' governments and is generally a low probability event for high-quality corporate bonds. A high-grade bond default is liable to coincide with or to follow an economic, financial, or political disturbance, which is typically accompanied by a shift in prevailing interest rates. Modified duration misses this correlation between interest rate moves and high-quality bullet bond defaults. Nonetheless, as long as credit risk is not a concern (e.g., deemed unlikely over the life of a bond) or is handled separately, modified duration can be an adequate measure of a bullet bond's interest rate sensitivity.

The situation is very different for bonds with optionality and for many kinds of ABSs. For instance, a callable bond is likely to be called by the issuer if interest rates decline, while the holder of a puttable one would probably exercise his right after a rate increase; both events would drastically alter the cash flows from the "bullet" pattern of a sequence of regular coupon payments topped with the principal repayment. Similarly, interest rates are central to the valuation of mortgage-backed securities, largely through the rate changes' radical effect on prepayments, as discussed in section 3.2. A measure that ignores changes in cash flows in response to interest rate moves would be meaningless—and, in fact, misleading—for these assets.

A linear rate sensitivity that incorporates the effect of discount curve shifts on cash flows is called *effective duration*. Given a valuation model, calculating effective duration is straightforward⁸:

$$\mathcal{D} \approx -\frac{1}{V} \frac{V(r_f(t) + \Delta r_{\parallel}) - V(r_f(t) - \Delta r_{\parallel})}{2\Delta r_{\parallel}} \quad (4.52)$$

where $V \equiv V(r_f(t))$ (i.e., calculated with unshifted discount curve) and the value for a shifted curve is calculated accounting for the full effect of the discount curve shift, including on the cash flows. Taking the limit of small $\Delta r_{\parallel}$, one arrives, analogously to the definition in equation (4.45), at

$$\mathcal{D} = -\frac{1}{V} \left. \frac{dV}{d\Delta r_{\parallel}} \right|_{\Delta r_{\parallel}=0}. \quad (4.53)$$

The difference with equation (4.45) is that now the calculation fully incorporates the effect of the interest rate change on the asset value, via both the discount curve and the cash flows. The mathematical expression of this property is that effective duration is minus scaled *total* derivative of asset value with respect to a parallel shift in the discount curve, a fact indicated by the total derivative symbol “d”; by contrast, modified duration is a *partial* derivative. We will discuss the practicalities of numerical evaluation of duration and other sensitivities in the next subsection.

Effective duration is similar to modified duration in the sense that both are exact sensitivities to parallel rate shifts; the new feature is the incorporation of the full effect of the rate shock. One can, therefore, use effective duration to calculate the asset value change in response to a small parallel shift $\Delta r_{\parallel}$ in the discount curve:

$$\frac{\Delta V}{V} \approx -\mathcal{D} \Delta r_{\parallel}. \quad (4.54)$$

This relationship extends equation (4.49) to reflect the rate shift’s impact on cash flows.

For many ABSs and bonds with optionality, modified and effective durations are vastly different. For some securities, most famously interest-only collateralized mortgage obligations (IOs), the two durations even have opposite signs: The effective duration of an IO is typically negative, as

⁸This is the most commonly used of a number of standard ways of calculating effective duration numerically. We will elaborate on the subject of numerical evaluation of sensitivities in the next subsection.

explained in section 2.2, whereas modified duration is always positive by definition.

When available, effective duration is usually preferable: It measures the true economic sensitivity to interest rate shocks (within the limitations of the valuation model, of course). Modified duration, however, can still be useful for bullet, especially high-grade, bonds. Its main advantage is the ease of computation: There is no need for a quantitative model; all that is required is the discount curve or even, since in practice the discounting is often done with the bond's yield to maturity, just a single observed yield. The consequent absence of model uncertainties is another argument in favor of using modified duration for bullet bonds.

In modern usage, except when applied to bullet bonds, "duration" usually means "effective duration" and, unless stated otherwise, we will adhere to this convention in the rest of this book.

PV01 and DV01

Duration is an asset's fractional value change in response to a rates shift. For a portfolio manager, it can be more important to know the dollar impact of a rates move. The commonly used measure is the so-called *dollar value of a basis point*, or DV01. This is the dollar amount by which the value of a portfolio or an asset changes in response to a one basis point *decline* in interest rates. Given duration, it is straightforward to obtain the DV01 from equation (4.45) by setting $\Delta r_{\parallel} = -0.0001$:

$$\text{DV01} = \mathcal{D} \times V \times 0.0001. \quad (4.55)$$

Since most assets' value increases when interest rates decline, a DV01 defined in this way is usually positive. Note, however, that sometimes the DV01 is defined with an opposite sign, that is, as an asset's value change caused by a one basis point *increase* in rates.

Outside the United States the quantity, defined in equation (4.55), is usually known as *present value of a basis point*, or PV01.

Convexity

When discussing the Greeks, we illustrated in Figure 4.1 the dramatic reduction in approximation error achieved by using a second-order sensitivity (gamma), in addition to a linear one. Similarly, the quality of the approximation in equation (4.54) can be improved by adding a term quadratic in rates change. The corresponding second-order sensitivity is called *convexity*

C and, apart from scaling with the asset's value, is similar in its definition to gamma, equation (4.22):

$$C = \frac{1}{V} \left. \frac{d^2 V}{d\Delta r_{\parallel}^2} \right|_{\Delta r_{\parallel}=0}. \quad (4.56)$$

We will usually use “effective” convexity, that is, the kind incorporating the full impact of a rate shift, and, therefore, as in the definition of effective duration, equation (4.53), we use the total, rather than partial, derivative.

We can now extend the approximation in equation (4.54) in a manner analogous to equation (4.23):

$$\frac{\Delta V}{V} \approx -D \Delta r_{\parallel} + \frac{1}{2} C \Delta r_{\parallel}^2. \quad (4.57)$$

As is the case with gamma, positive convexity is generally desirable: It enhances the upside and mitigates the downside. Most bonds' and asset-backed securities' convexity is usually positive. An example of an instrument that can exhibit negative convexity is a callable bond, as described in section 2.2. Extending the parallel with gamma, an excessive convexity, even positive, can hinder the maintenance of a duration-neutral portfolio by generating duration in response to rate moves. To avoid frequent rebalancing, managers can deliberately keep the convexity of their duration-neutral portfolios low.

Key Rate Measures

The rate sensitivity measures we have described thus far quantify an asset's response to *parallel shifts* in the discount curve. These quantities do not differentiate between a 5 basis points spike in short-term three-month rates and an equal increase in long-term 10-year yields. This approach is too crude for many bonds and asset-backed securities whose value often depends in a complex and subtle way on the term structure of interest rates. We presented a simple example in section 2.2 when discussing yield curve risk: Two sample portfolios had equal durations but very different exposures to the yield curve term structure. While these portfolios would respond similarly to a parallel shift, a (more realistic) change in the curve's *shape* would in general induce different responses.

In order to manage the full rates exposure, a manager needs to quantify his portfolio's sensitivity to the shape of the discount curve. One commonly used approach is to shift a segment of the curve in the vicinity of one of a set of predetermined *key maturities* while keeping the remaining rates

unchanged. The resulting *key rate measures* mirror the sensitivities to parallel shifts and are known as *key rate duration*, *key rate DV01* (or *PV01*), and *key rate convexity*.

It is easiest to understand these quantities on the example of a bullet bond assuming that cash flows are unaffected by rates changes (i.e., ignoring credit risk). In this case the expressions for key rate measures, which are *modified* in the sense of equation (4.45), can be calculated exactly and have a particularly simple form.

Let us consider a bullet bond whose value, using continuously compounded discount rates, is

$$V = \sum_{i=1}^N C(t_i) e^{-r_f(t_i)t_i} \equiv \sum_{i=1}^N C_i e^{-r_i t_i} \quad (4.58)$$

where for convenience we introduced the notation $C_i \equiv C(t_i)$ and $r_i \equiv r_f(t_i)$, and t_i is time to cash flow C_i . For this bond it is natural to consider cash flow dates t_i as key maturities and, accordingly, rates r_i as key rates. Then the i th key rate duration (for each of $i = 1, \dots, N$) is minus-scaled sensitivity of V to r_i , keeping all other rates constant:

$$\mathcal{D}_i = -\frac{1}{V} \frac{\partial V}{\partial r_i} = \frac{1}{V} t_i C_i e^{-r_i t_i} \quad (4.59)$$

(a mathematically savvy reader would have recognized the analogy between key rate measures and functional derivatives). Defining key rate DV01_{*i*} by analogy with equation (4.55), we obtain

$$\text{DV01}_i = \mathcal{D}_i \times V \times 0.0001 \quad (4.60)$$

Similarly, key rate convexities are

$$C_{ii} = \frac{1}{V} \frac{\partial^2 V}{\partial r_i^2} = \frac{1}{V} t_i^2 C_i e^{-r_i t_i}. \quad (4.61)$$

Note that convexity has a double index. The reason is that, as do gammas, key rate convexities form a symmetric matrix, and a general definition of *effective* key rate convexities should be

$$C_{ij} = \frac{1}{V} \frac{d^2 V}{dr_i dr_j} \quad (4.62)$$

(compare this to the definition of gammas in equation (4.24)). It so happens that for a bullet bond with cash flows independent of rates, the off-diagonal elements (those with $i \neq j$) vanish. For a more complex asset off-diagonal convexities are not identically zero because cash flows depend on multiple key rates: A five-year rate can affect the seventh-year cash flow. We leave it as an exercise for the reader to show that for bullet bonds, key rate measures sum to their modified counterparts:

$$\sum_{i=1}^N \mathcal{D}_i = \mathcal{D}_{\text{mod}}, \quad \sum_{i=1}^N \text{DV01}_i = \text{DV01}_{\text{mod}}, \quad \sum_{i=1}^N C_{ii} = C_{\text{mod}} \quad (4.63)$$

(DV01_{mod} is calculated as in equation (4.55) but using $\mathcal{D}_{\text{mod}}$ and modified convexity is calculated as $C_{\text{mod}} = V^{-1} \partial^2 V / \partial \Delta r_{\parallel}^2$, i.e., analogously to equation (4.56) but with cash flows kept unchanged).

Extending equation (4.57), a fractional change in asset valuation, resulting from a curve shift, can be expressed in terms of effective key rate durations and convexities:

$$\frac{\Delta V}{V} \approx - \sum_i \mathcal{D}_i \Delta r_i + \frac{1}{2} \sum_{ij} C_{ij} \Delta r_i \Delta r_j \quad (4.64)$$

where the summation extends over all key rates: $i, j = 1, \dots, N$. For a bullet bond, it is easy to show, using equation 4.63 and the observation that $C_{ij} = 0$ for $i \neq j$, that in the case of a parallel shift (i.e., $\Delta r_i = \Delta r_{\parallel}$ for all i) this expression correctly reduces to equation 4.57.

For a bullet bond, the choice of coupon dates as key maturities is natural, and the resultant key rate decoupling (each term in equation (5.58) involves a single key rate) makes calculations easy. For asset-backed securities or bonds with embedded options, the selection of key maturities is less obvious: These assets' value can depend on every point on the continuous discount curve, with each cash flow potentially affected by rates on all the preceding dates, in addition to the payment date itself. While one could formally calculate key rate measures for every day of the remaining 25 years of an MBS's life, this would be a pointless exercise since any useful information would be drowned in the vast output and subsumed by model errors. Instead, the usual practice is to select a limited number of key maturities that representatively cover the span of the discount curve (which can extend, depending on the asset, to 30 years and beyond). A sensible choice might be to use the terms of the reference rates, introduced in section 3.3, as key maturities. This selection might help to answer another important question related to key measure calculation: How should one *deform* the curve when a single key rate is shifted?

That is, if, by way of an example, one increases the five-year rate by one basis point and the nearest key maturities are three and 10 years, how should the region of the discount curve between three and 10 years change? An obvious solution, aided by the selection of reference rate terms as key maturities, might appear to be to rerun the curve interpolation model with the shifted reference rate: This would ensure consistent methodology and produce a shifted curve with the same degree of smoothness as the original one. One should, however, beware of possible undesirable side effects. As follows from our discussion in section 3.3, some interpolation techniques can produce discount curves that differ *outside* the adjacent key maturities (three and 10 years in our example), and curves produced by certain methodologies may not even pass through the reference rates. For this reason practitioners often opt for a simple piecewise linear shape for the *key rate deformation* (i.e., the difference between the shifted and the unshifted discount curves), illustrated in Figure 4.2. To shock the five-year rate, for example, one increments the curve linearly from the previous key rate (three years) to the five-year point and then decrements it, linearly, to the next key rate, 10 years. Note that the discount curve remains unchanged—the deformation in Figure 4.2 equals zero—beyond the key rates (three and 10 years), neighboring the shifted one (five years). Shifts in outermost key rates, one




FIGURE 4.2 Piecewise linear key rate deformations for one basis point shifts in one, five, and 10-year key rates. The key rates are one, three, five, and 10 years.

and 10 years in this example, result in persistent deformations in the outlying regions. Thus, the impact of a one-basis point shock to the 10-year rate is linear growth to its left, starting at the adjacent key rate of five years, and a constant one basis point shift to its right, for as long as the discount curve extends. Similarly, a shock at one year uniformly elevates the region from the shortest (overnight) rates to the one-year point; the deformation recedes linearly to zero between one year and the next key rate, three years in this example.

The piecewise linear approach has an additional advantage of allowing the flexibility of matching the key maturities to the specifics of the asset, regardless of the reference rates.

Other Discount Curve Sensitivity Measures

Depending on the objective, other measures of sensitivity to discount curve shape are also used. One approach is to decompose the curve move into a parallel shift and a sequence of standardized de-measured perturbations of increasing complexity (curve steepening, butterfly, etc.) and to calculate asset sensitivities to these perturbations. Another is to approximate the discount curve with a polynomial, typically cubic, in times to maturity and to calculate responses to changes in linear, quadratic, and cubic coefficients. See Tuckman (2002) for details.

4.6 NUMERICAL EVALUATION OF SENSITIVITIES

We defined sensitivities as partial or total derivatives of an asset's value with respect to various continuous factors. Now we discuss the practical aspects of their numerical calculation using a valuation model.

Let us consider an asset, whose value V we can calculate by using a financial model and let us focus for convenience on a single market factor f . We wish to calculate V 's sensitivity χ , defined in equation (4.5) as the derivative of V with respect to f . Since the derivative is the limit of the ratio $(V(f+u) - V(f))/u$ as u tends to zero, one might be tempted to use this ratio for a small but finite u as a suitable approximation:

$$\chi \approx \frac{V(f+u) - V(f)}{u}. \quad (4.65)$$

This estimate, however, is almost guaranteed to be very inaccurate. One can show (Press et al., 2002) that if one can compute V with a fractional accuracy of ε , then equation (4.65) *at best* has a fractional accuracy of

$$\frac{\text{Error}[\chi]}{\chi} \approx \sqrt{\varepsilon} \quad (4.66)$$

achieved when the shift u has an optimal value of

$$u_{\text{opt}} \approx \sqrt{\varepsilon} f_{\text{curv}}, \quad f_{\text{curv}} = \left| \frac{V}{\Gamma} \right|^{1/2}. \quad (4.67)$$

Here f_{curv} is the so-called *curvature scale* of $V(f)$, namely, the typical magnitude of shift Δf over which the quadratic term in expansion (4.23) becomes comparable to $V(f)$. The presence of Γ in equation (4.67) appears to lead to a circular argument: One needs a second-order sensitivity (Γ) in order to calculate optimally a first-order one. In practice, one can use a “quick and dirty” estimate of Γ , obtained with a reasonable non-optimal argument shift, in the calculation of f_{curv} . In fact, often one can simply set $f_{\text{curv}} = f$ (unless, of course, near $f = 0$).

The implications of equation (4.66) are quite remarkable. It says that if a model uncertainty in valuing an asset is 1%, then the best accuracy one can hope to achieve in calculating sensitivities is 10%. For a more realistic (but often still optimistic) model uncertainty of 5%, the sensitivity error would exceed 20%. A model error of 10% would be blown in the best case to almost one third.

While error amplification when calculating sensitivities is unavoidable, it can be mitigated by replacing approximation (4.65) with a symmetric one

$$\chi \approx \frac{V(f+u) - V(f-u)}{2u}. \quad (4.68)$$

Without going into details, this modification replaces a truncation error linear in u (approximately Γ) with a smaller quadratic one ($u^2 \times \text{Speed}$; see equation (4.39)). The optimal shift now is

$$u_{\text{opt}} \approx \varepsilon^{1/3} \left| \frac{V}{\text{Speed}} \right|^{1/3}, \quad (4.69)$$

leading to a fractional error of

$$\frac{\text{Error}[\chi]}{\chi} \approx \varepsilon^{2/3}. \quad (4.70)$$

This is smaller than equation (4.66) by a factor of $\varepsilon^{2/3}/\varepsilon^{1/2} = \varepsilon^{1/6}$. For $\varepsilon = 5\%$, the improved best fractional accuracy is 14% and for $\varepsilon = 10\%$ it is 22%; these numbers should be compared with 22% and 32%, respectively, for the estimate in equation (4.65). Still, the reduced errors greatly exceed the model uncertainty ε , and analysts and managers should be mindful of the error magnitude and its implications.

Errors are amplified again when calculating higher order sensitivities. Roughly, the lower error bound (4.70) now becomes the input, so that for a numerically calculated gamma, the fractional accuracy is not better than

$$\frac{\text{Error}[\Gamma]}{\Gamma} \approx (\varepsilon^{2/3})^{2/3} = \varepsilon^{4/9}, \quad (4.71)$$

substantially greater than bound (4.70). Thus, for a model uncertainty of 5%, the gamma fractional accuracy is at best 26%, almost twice that for χ .

Errors will become progressively amplified for each successive sensitivity order. This limits practical usefulness of numerically computed higher-order sensitivities.

There exist techniques for achieving accuracy better than the bound in equation (4.70). One of the more conceptually appealing is *Richardson extrapolation* in which one calculates a numerical estimate $\chi(u)$ for a range of decreasing values of u and extrapolates the result to $u \rightarrow 0$ (Press et al., 2002). This and other enhanced accuracy methods require multiple evaluations of V and may be impractically expensive computationally.

4.7 PERFORMANCE ATTRIBUTION AND COMPLETENESS TEST

Performance Attribution

In the preceding section, we defined factor sensitivities and described the process of calculating them using a valuation model. We also showed how to convert sensitivities to equivalent variation investments, a crucial tool in hedge sizing. Another important use of sensitivities is in *performance attribution*, that is, the breaking up of an asset value change over a time period into parts, caused by moves in distinct factors. If a model returns an asset value of $V(t_0)$ at time t_0 and a value of $V(t_1)$ at a later time t_1 , then one must be able to explain the profit or loss $\Delta V \equiv V(t_1) - V(t_0)$ in terms of the factor moves during the interval $\Delta t = t_1 - t_0$. At the linear level, performance attribution is achieved by equation (4.7) if Δf_{ks} are understood

to be changes in factor levels over the period Δt . Equation (4.25) extends attribution to the quadratic level. (These expressions do not formally include the effect of interest rate moves. We will address these effects later in this section.)

Linear attribution unambiguously assigns each term in the decomposition: value change, attributable to the k th factor, is

$$\Delta V_k = \chi_k \Delta f_k \quad (\text{linear}). \quad (4.72)$$

At the quadratic level the assignment is less clear-cut since cross gammas entangle different factors. As a practical approach, an analyst might combine the diagonal quadratic (and higher, if used) terms with the linear ones and group all the off-diagonal contributions together as a measure of the effect of cross-factor interaction:

$$\begin{aligned} \Delta V_k &= \chi_k \Delta f_k + \frac{1}{2} \Gamma_{kk} \Delta f_k^2, \\ \Delta V_{\text{cross}} &= \frac{1}{2} \sum_{k \neq k'} \Gamma_{kk'} \Delta f_k \Delta f_{k'} \quad (\text{quadratic}). \end{aligned} \quad (4.73)$$

In real-life applications, depending on the nature of the asset, the details of the model, and the objective of the calculation, the manager may compute on a routine basis only the first-order sensitivities, or, say, a combination of all first-order sensitivities and a selection of second- (and, sometimes, higher-) order ones. When using these results for performance attribution, one should take care not to omit important terms. It would be negligent, for instance, to calculate and include in the attribution analysis the gamma to the Euro STOXX 50 but not to the S&P 500 if the latter's gamma happens to be comparable or greater in magnitude (typical changes in the two indices would be similar in size).

The sum of all the linear and nonlinear attributions,

$$\Delta V_{\text{attrib}} \equiv \sum_k \Delta V_k + \Delta V_{\text{cross}}, \quad (4.74)$$

would in practice differ from the model value change ΔV . The unattributed portion of the P/L,

$$\Delta V_{\text{unattrib}} = \Delta V - \Delta V_{\text{attrib}}, \quad (4.75)$$

reflects largely the effects of higher-order nonlinearities not included explicitly in the attribution process and should in general be small compared to ΔV . An exception to this rule is when ΔV itself is expected to be small, for example, if factor changes Δf_k are close to zero (e.g., smaller in magnitude than their respective historical standard deviations). Barring such particular circumstances, the smallness of $\Delta V_{\text{unattrib}}$ is a measure of the degree of *self-consistency* of the model, which produces both the value change and factor sensitivities. Absent implementational model errors, a persistently large unattributed portion of ΔV may indicate that some of the omitted terms, say, a particular gamma, are in fact important and should be incorporated in the attribution process.

Sensitivities used in attribution analysis can in principle be calculated at either time t_0 or t_1 (or, for that matter, at any time between t_0 and t_1). The standard, especially when using sensitivities of order higher than first, is to use t_0 . In addition to conforming to the standard (but by no means unique) definition of mathematical derivatives, this approach avoids spurious, and potentially dangerous, misattribution following non-market events. To illustrate the danger, consider a portfolio whose value at t_0 was $V(t_0)$ and that received a new investment equal to $V(t_0)$ just before t_1 . Sensitivities, calculated at time t_1 , would be roughly twice those at t_0 , but simply because of the doubling of the portfolio size, not as a result of market moves or non-market factor evolution. An analyst might decide to use the sensitivities at t_1 for convenience or the average of values at t_0 and t_1 , believing, not without justification (see next paragraph), that such averaging would improve accuracy. Should he do so, however, he might be puzzled by a very large $\Delta V_{\text{unattrib}}$, even after accounting for cash inflows. The reason is that this approach would greatly overestimate sensitivities as they applied for most of the period. Using sensitivities, calculated at t_0 , and accounting for the inflows should result in a small unattributed P/L.

Having said this, one can make a legitimate case for averaging sensitivities between the end points of the calculation period: The averaging allows one effectively to increment by one the order of the sensitivity calculation with little additional effort. In order to illustrate this effect, let us focus on a single factor, a price p , with a corresponding delta δ and gamma Γ . Suppose an analyst uses the average delta, $\frac{1}{2}(\delta(t_0) + \delta(t_1))$ to calculate the portion of the P/L, attributable to a change in price Δp . Since the change in the delta itself is determined by gamma, according to equation (4.26), the P/L attribution is

$$\begin{aligned}\Delta V[\text{attributed to price change } \Delta p] &= \frac{1}{2}(\delta(t_0) + \delta(t_1))\Delta p \\ &\approx \frac{1}{2}(\delta(t_0) + [\delta(t_0) + \Gamma(t_0)\Delta p])\Delta p = \delta(t_0)\Delta p + \frac{1}{2}\Gamma(t_0)\Delta p^2. \quad (4.76)\end{aligned}$$

In other words, the analyst would calculate the attribution to second order (ignoring cross terms) without explicitly evaluating the gamma. When managing a complex asset, whose value is expensive to calculate, this trick could be a practical means to improve the quality of attribution without additional computation. The analyst, however, should be aware of the potential pitfall in case of a non-market value change, such as cash inflows or outflows.

Attribution to Interest Rate Moves Thus far in this section we have phrased the discussion in terms of “equity-style” sensitivities. The extension of the attribution process to interest rates is straightforward: One should add to the right-hand side of equation (4.72) and of ΔV_k in equation (4.73) the valuation changes resulting from interest rate moves according to, respectively, equations (4.54) and (4.57). If key rate analysis is used, then one needs to add the linear (duration-induced) terms in the ΔV , calculated from equation (4.64), to equation (4.72), and ΔV_k in equation (4.73). The diagonal quadratic terms of equation (4.64), that is, those with $i = j$, will also contribute to ΔV_k , while the off-diagonal cross terms should be added to ΔV_{cross} .

Completeness Test

We noted earlier in this section that a large unattributed P/L $\Delta V_{\text{unattrib}}$ may signal a deficiency in the attribution procedure. Additionally, time series of factor-attributed and unattributed P/Ls can help assess the soundness of the model, specifically, the degree to which the model’s set of factors captures all the significant market and non-market influences. If the set of factors is adequately selected, then $\Delta V_{\text{unattrib}}$ should be *contentless noise*, in particular, it should be uncorrelated with the factors or with ΔV_{attrib} .

A useful practical test is to calculate the correlation coefficient⁹ ρ_{complete} of the empirical time series of $\Delta V_{\text{attrib}}(t_i)$ and $\Delta V_{\text{unattrib}}(t_i)$; these series could be daily, weekly, or monthly, depending on the frequency with which the model is run. For a *complete model*, whose factors comprise all the important influences on the asset value, ρ_{complete} should equal nought. A statistically significant non-zero correlation between the attributed and unattributed portions of the P/L suggests that an important value driver is missing from the model. This test becomes a particularly powerful probe of the quality of the model’s underlying assumptions if the valued asset is traded in sufficiently

⁹We define correlation coefficient and describe statistical tests in Appendix B.

liquid markets and valuations $V(t_i)$ are empirically observed prices (while the sensitivities and, consequently, $\Delta V_{\text{attrib}}(t_i)$ s are model-derived).

Readers can find a practical model completeness investigation of commercial risk and attribution systems in Bychuk, Kurinets, and Schadt (2002). The authors of this study performed out-of-sample analysis using actual returns of a group of mutual funds.

CHAPTER 5

Quantifying Portfolio Risks

... to pursue mathematical analysis while at the same time turning one's back on its applications and on intuition is to condemn it to hopeless atrophy.

Richard Courant, *Differential and Integral Calculus*

5.1 THE NATURE OF RISK

Risk is ultimately an impalpable and, to a large degree, subjective concept, reflecting a multitude of tangible and intangible factors, perceptions, and sentiments, some measurable, some observable, some quantifiable, and some identifiable, but many more not. Risk management and investment activity in general strive to accomplish an epistemologically dubious task of predicting, and anticipating, the future using the past as a guide.

Of course, projecting past behavior into the future as a portfolio management approach is unavoidable, if only because there are few alternatives. This projecting relies on the persistence of the statistical properties of capital market observables (prices, trading volumes, interest rates, foreign exchange rates), both individually, as captured, for example, by the probability distribution of past returns and its cumulants, and in relation to each other, as measured, for instance, by pair-wise correlations of returns. It also assumes the similarity of past and future responses of market and economic indicators to discontinuous events, such as government and central bank policy announcements, legal actions, or terrorist attacks. In reality, both statistical properties and responses to shocks usually change in volatile or dysfunctional market conditions.

An illustrative example is the unexpected announcement by the U.S. Federal Reserve of its plans to purchase \$300 billion longer-term Treasury notes and bonds, as well as \$1 trillion mortgage-backed securities, on March

18, 2009. Usually, U.S. stock prices and Treasury yields move in unison: When equities rally and the sentiment is positive, funds are flowing from the safety of government securities to the risky stock markets, Treasury prices fall, and yields rise; when investors are fearful, they “flee to quality” into Treasuries, and both stocks and yields fall. One quantitative expression of this relationship is the correlation coefficient of daily returns on the S&P 500 index and daily changes in 10-year Treasury yields; between March 2007 and February 2009, this correlation coefficient was +0.5.

On the day of the announcement, the S&P 500 rose 2.09% while the 10-year yield *fell* 47.4 basis points, 5.8 standard deviations of daily changes over the same period, March 2007 to February 2009, and the biggest one-day fall in 20 years. In fact, the equity index rose 1.1% and the yield fell 32 basis points within four minutes of the announcement, made at 2:15 pm New York time, as illustrated in Figure 5.1. At its lowest, at approximately 2:34 pm, the yield declined more than 54 basis points from the previous day’s close, the biggest intraday drop since 1962.

Another example is the well-documented fact that correlations between stock returns become large positive across the spectrum in declining markets, a phenomenon also known as *contagion* (see, for instance, a study by Bekaert, Harvey, and Ng, 2005). This is a sophisticated way of saying that everything tanks together, and portfolio diversification fails precisely when one needs it the most. The case is well illustrated by Chicago Board Options Exchange (CBOE) S&P 500 Implied Correlation Indices. These are representative correlations, inferred from the observed prices of options on the S&P 500 and on its components; the reader can find the details of the methodology at CBOE’s website (cboe.com). Figure 5.2 shows rebased values of one of these indices, ICJ, derived from options expiring in January 2010. Its comparison with the S&P 500 Total Return Index reveals that correlations among S&P 500 constituents shot up when the returns slumped, peaking (with a value of 75%) about two weeks after the stock market’s nadir. The correlation coefficient of S&P 500 daily total returns and daily changes in the ICJ was –31% during the period shown, thus confirming the visual observation.¹

This increased correlation in bear markets, while being detrimental to diversification, can, in fact, aid the task of *hedging*: An ideal hedging instrument would be perfectly negatively correlated with the hedged asset, at least in declining markets. Thus, the short side of a futures contract on

¹The stated value of the correlation coefficient has a high statistical significance, Student’s *P*-value being 10^{-12} . See Appendix B for a discussion of statistical significance measures.




FIGURE 5.1 Market reaction to the March 18, 2009 U.S. Federal Reserve announcement of the plans to purchase Treasuries and mortgage-backed securities. The graphs show an abrupt increase in S&P 500 (top) and an even more dramatic drop in the U.S. Generic Government 10-Year Yield Index (bottom) immediately following the Federal Open Market Committee (FOMC) statement release at 2:15 pm. These movements contravene the historical tendency for the two indices to move in tandem. Used with permission of Bloomberg L.P. and its subsidiaries. Copyright© 2010. All rights reserved.




FIGURE 5.2 Rebased CBOE S&P 500 Implied Correlation and S&P 500 Total Return indices. The graph shows the period between November 19, 2007, and November, 23, 2009, when CBOE calculated an implied correlation index (ICJ) that was using options expiring in January 2010.

Source: Bloomberg Finance LP.

the S&P 500 might normally track a particular portfolio of U.S. stocks imperfectly but is likely to provide a surprisingly good hedge when the stock market drops.

Historical relationships usually describe a portfolio's behavior reasonably well in *normal* circumstances, that is, under the conditions that obtained *on average* during the considered historical period, typically quite extended. History, however, does not repeat itself, even if it rhymes, and relationships underlying hedging or diversification decisions break down when the financial markets enter a period of turbulence.

A market practitioner, nonetheless, needs tangible quantitative measures of dangers lurking in his portfolio in order to assess, budget, and control his market risks. In the following sections, we describe the commonly used quantitative risk measures while urging the reader to be mindful of their limitations and the dangers of their uncritical application.

A portfolio manager can to a certain extent prepare for the unexpected by approaching the problem from the opposite end. Instead of asking, "What is likely to happen?" he may ask, "What would have to happen to inflict unacceptable losses on the portfolio?" and "How would the portfolio fare in a 'perfect storm' when all conceivable adverse events occur simultaneously?"

In other words, the manager wants to study how a portfolio would behave under an extreme stress, or conduct a *stress test*. Stress tests, discussed in section 4.3, still can rely on historical relationships via the asset model assumptions, but to a lesser degree and in an indirect fashion. Armed with the results of a comprehensive set of stress tests, the manager can take considered steps to mitigate the effects of unforeseen extreme events.

The most important rule is not to rely excessively on any one approach but rather to combine historical analysis with stress testing and *scenario analysis*. The latter seeks to identify those sets of circumstances that will result in a specified loss to the portfolio.

A successful manager combines factual data and quantitative analysis with a deep understanding of financial markets and the broader economy, of common sense, and *intuition*, understood as a manifestation of experience and acumen. It is important to have a wide-ranging view of all facets of capital markets and economic activity, as well as human society as a whole, and how they interact and influence each other.

We particularly emphasize the notion of risk as a *concept* as opposed to a *number* or collection of numbers. In common practice “risk” is often used interchangeably with “risk measure,” especially one specific, and dangerously limited, measure, the standard deviation of past returns. While this shorthand may be useful and should not be stigmatized as long as it is unambiguous, a market practitioner should be at all times acutely aware that any numerical measures present only certain aspects and a simplified view of a complex underlying entity, the tail, leg, trunk, ear, or tusk of the proverbial elephant examined by a group of blind men.

5.2 STANDARD RISK MEASURES

In Chapter 4, we described how to use a valuation model to quantify an asset’s exposures and calculate its sensitivities to market and other factors. Exposures and sensitivities relate an asset’s properties to extraneous factors: They provide *extrinsic* measures of risk. In order to produce response matrices, scenario grids, and factor sensitivities that quantify extrinsic risk, we used the model *as a whole* (i.e., as a “black box”), applying its computational engine $\mathcal{M}$, introduced in equation (3.2), to a range of purposefully selected sets of input parameters $\mathcal{I}_q$. We did not need to know the details of the model or its implementation.

In this and the next sections we introduce measures of asset risk that are free from explicit reference to external factors. These *intrinsic* risk measures come in pairs: Each measure can be *realized* and *predicted*. The realized versions are calculated from an asset’s past performance (or, in the case of relative risk measures, from the past performance of two or more assets) and

are completely independent of any quantitative model. Predicted intrinsic risk measures must be generated by a risk model and, unlike sensitivities, they cannot be obtained by rerunning a valuation model as a whole.

Of course, a realized (i.e., historical) performance statistic is in itself a predictor of an asset's likely future behavior and, unlike a forecast, it has the advantage of being free from model risk. A realized statistic may not, however, reflect all the latest information, it may carry the baggage of no-longer relevant events that occurred months or even years ago, and its value would depend on subjective choices, such as the evaluation period, data frequency, and, if applicable, the rate of historical decay. In addition, the asset may simply not have performance history or only have a history of length, insufficient to produce statistically significant results (this can be particularly an issue with tail risk measures discussed in the next section).

A model would calculate predicted statistics using, possibly implicitly, the probability distribution of an asset's returns (see discussion in section 3.3). Of course, this probability distribution reflects past observations and, therefore, forecast statistics share some of the disadvantages with realized ones. A forecast, however, can be more reliable than a simple historical measurement of the desired statistic because good models make projection in an economical and robust way, for instance, by means of stochastic extrapolation, discussed at the end of section 3.3. A manager should, where possible, utilize both the observed and the model-projected statistics as complementary measures.

There are numerous commercially available and proprietary risk models that predict the measures we define below for various asset classes.² We will provide conceptual descriptions and statistical definitions (i.e., definitions stated in terms of historical performance) and discuss each measure's significance, implications, and limitations. We expect practitioners to have access to risk systems capable of predicting their assets' and portfolios' risk characteristics.

Absolute Risk Measures

Volatility: Standard Deviation and Variance of Returns Standard deviation of returns, frequently referred to simply as "volatility," is the most commonly used risk measure. In fact, "risk" has often come to mean, rather loosely, volatility of returns. This statistic's popularity derives from the ease of calculation and interpretation, as well as the useful analytical properties of its square, the variance.

²For an empirical study of the quality of commercial risk systems' predictions, the reader can consult Bychuk, Kurinets, and Schadt (2002) and references therein.

Variance and standard deviation both measure the tendency of an asset's returns to vary. All else equal, an investor prefers low variability of returns: 5% annually with near certainty (i.e., with low volatility) is better than the same expected return of 5% but with a high likelihood of receiving less than say -5% (as well as more than $+15\%$), that is, with a volatility of some 10% or higher; for the same expected return, investors prefer to assume lower risk, that is, less uncertainty.

The variance is a *cumulant* of the probability distribution of returns, a characteristic, introduced in section 3.3, and as such is additive across time. It is not, however, directly comparable to returns as it has the dimensions of “return squared” (e.g., of “percent squared” if the returns are expressed in percentage points). Standard deviation, the square root of variance, is, by contrast, comparable to returns, and a statement “the expected return is 5% with a standard deviation of 10%” gives an immediate sense of the likely range of realized outcomes.

Given a time series of historical prices or valuations $\{p(t_i)\}_{i=0}^N$, the returns variance is defined in terms of logarithmic returns, introduced in footnote 16 of Chapter 3:

$$\sigma^2 = \mathcal{A} \times \frac{1}{N-1} \sum_{i=1}^N (r_i - \bar{r})^2$$

$$\text{where} \quad r_i = \ln \frac{p(t_i)}{p(t_{i-1})} \quad \text{and} \quad \bar{r} = \frac{1}{N} \sum_{i=1}^N r_i. \quad (5.1)$$

The *annualization factor* $\mathcal{A}$ ensures that the statistic is expressed in terms of annual returns. The value of $\mathcal{A}$ depends on the frequency of prices $p(t_i)$: For monthly observations, the annualization factor is 12; for weekly ones it is 52; and for daily data, $\mathcal{A}$ is the number of business days in the year, which, of course, varies and is usually set to a number between 250 and 260. Standard deviation σ is the square root of the variance.

One should use $N-1$ in the definition of σ^2 if the mean $\bar{r}$ is estimated from the same data set, as in equation (5.1). If the mean is known *a priori* (e.g., if an analyst postulates that the mean return is zero), then one should replace the $N-1$ with N . The reason for this can be found in any standard statistics book. To quote from Press et al. (2002), “if the difference between N and $N-1$ ever matters to you, then you are probably up to no good anyway—e.g., trying to substantiate a questionable hypothesis with marginal data.”

If the financial instrument pays distributions (dividends, coupons, etc.) and the analyst is interested in volatility of total returns, then $\{p(t_i)\}_{i=0}^N$

should be a time series of *total return levels*—that is, a period change in p should equal the sum of the price change and period distributions. Using total rather than price returns typically makes little difference to historical standard deviations of stocks and stock indices, but doing so can be decisive for bonds and other fixed income securities.

Constructing total return level time series in the way just described assumes reinvesting distributions in the asset. This may not be possible with illiquid instruments, for example, certain corporate and municipal bonds and asset-backed securities (even when these instruments are available for purchase high transaction costs can make reinvestment impractical). In addition, the procedure assumes that the reinvestment occurs at the end of each period, while in reality distributions can arrive at any time, and the price evolution between the distribution and the period end is ignored. The resulting distortion would be particularly pronounced when low frequency data (quarterly, monthly, or weekly) are used. Because of these complications, standard deviation of total returns may not be a good measure of the riskiness of fixed income securities—interest rate volatility in conjunction with duration and convexity would often be more appropriate.

Semivariance and Semideviation Variance and standard deviation equally penalize returns below and above the mean. An investor would be wary of an asset that can generate a large loss (and he would require additional expected return for this downside risk) but would generally be less concerned with atypically large positive returns. In fact, actively (and well) managed funds tend to exhibit more upside than downside large moves. A statistic that punishes underperformance but is indifferent to higher than expected upside is *semivariance*:

$$\sigma_-^2 = \mathcal{A} \times \frac{1}{N-1} \sum_{r_i < \bar{r}} (r_i - \bar{r})^2 \quad (5.2)$$

where r_i s and $\bar{r}$ are defined in equation (5.1), and the summation extends only over below-average returns. As in the case of variance, a more intuitive and practically convenient quantity is the square root σ_- of semivariance known as *semideviation*.

Semideviation penalizes all returns that are below period average. An investor may instead be concerned that the rate of return exceed a minimum acceptable level. For instance, he may wish to receive at least the risk-free rate r_f , or he may require that an investment compensate him for bearing the risk by returning an adequate risk premium in addition to r_f . A pension fund supervisor may determine that he needs at least 5% long-term *per*

annum return in order to meet projected fund liabilities. *Target semivariance* penalizes performance below a pre-determined *target return* r_{target} :

$$\sigma_{r_{\text{target}}}^2 = \mathcal{A} \times \frac{1}{N-1} \sum_{r_i < b} (r_i - b)^2, \quad b = r_{\text{target}}/\mathcal{A}. \quad (5.3)$$

Here we assume r_{target} to be annualized, b is the period target return, and all returns are logarithmic. *Target semideviation* $\sigma_{r_{\text{target}}}$ measures the risk that the asset or portfolio will fail to meet the minimum return requirement.

Given the apparent advantage over standard deviation of semideviation and target semideviation as a measure of downside risk, one might be surprised that they are not used more widely. The main reason is the assumption of log-normality of price moves that underlies most risk systems. Indeed, under the resulting normality of logarithmic returns, the use of semivariance is conceptually equivalent to using variance and would lead to the same investment decisions: The semivariance equals exactly half the variance, and a target semivariance is a predetermined fraction of the variance, calculated from the expected return and r_{target} . In fact, there is no benefit in using semivariance for any symmetric distribution of logarithmic returns. The situation would be different for a model using asymmetric probability densities. As we argued in section 3.3, to be meaningful, this would have to be an algebraically long-tailed Lévy distribution.

Semideviation and semivariance can be useful in the analysis of *historical* performance by providing comparable measures of various securities' downside risk or by comparing downside risk of the same asset during different time periods.

Sharpe Ratio The Sharpe ratio (Sharpe, 1994) is probably the second most familiar quantitative measure after standard deviation. Also known as *risk-adjusted return* and *return per unit of risk*, it is defined as return in excess of the risk-free rate divided by the standard deviation of returns:

$$\text{Sh} = \frac{r - r_f}{\sigma} \quad (5.4)$$

(all the quantities are annualized).

The Sharpe ratio penalizes assets, delivering the same return with higher volatility. It is comparable across various risky assets regardless of their nature, is straightforward to calculate historically, and can be easily generated by most risk systems.

It does, however, have shortcomings. One is the ambiguity in inputs, especially the standard deviation: The precise value of a historically calculated σ would depend on the length of the sample period and the sampling frequency. In addition, the choice of r_f may in practice be equivocal, as discussed in section 3.3. But the Sharpe ratio's most important limitation is that, just as the standard deviation, it penalizes positive returns. A related peculiarity is that, used uncritically, the ratio would favor investments with tiny excess returns and an even tinier volatility, such as some money market funds, over an asset returning, say, 5% above r_f with a volatility of 6%—most investors would prefer the latter.

By using only the expected return and its standard deviation the Sharpe ratio implicitly assumes that these two moments completely determine the returns' distribution, that is, that the (logarithmic) returns are normally distributed. The ratio, therefore, cannot capture—and penalize—tail risk.

Sortino Ratio The Sortino ratio (Sortino and Price, 1994) modifies the concept of the Sharpe ratio by measuring return per unit of *downside* risk:

$$So = \frac{r - r_{\text{target}}}{\sigma_{r_{\text{target}}}}. \quad (5.5)$$

This ratio can prove particularly useful in *ex post* evaluation of realized returns, especially when ranking money managers or comparing asset classes. As a forward-looking measure, however, it is subject to the same limitation as target semivariance: The use of the Sortino ratio is conceptually equivalent to the use of the Sharpe ratio when the underlying model assumes normal—or any symmetric—distribution of logarithmic returns (it is straightforward to generalize the Sharpe ratio to use a target return rate in lieu of r_f).

Treynor Ratio The Treynor ratio equals excess return per unit of *systematic risk* as measured by the security's or portfolio's *beta*, the slope coefficient β in a regression with respect to a broad market index:

$$Tr = \frac{r - r_f}{\beta}. \quad (5.6)$$

The Treynor ratio, most relevant for long-only equity investments, embodies the capital asset pricing model's fundamental premise that a stock's risk can be decomposed into *systematic* and *idiosyncratic* components (the latter also known as *unsystematic* risk). According to CAPM, idiosyncratic risk, quantified by the intercept—or alpha—in the regression of the stock's excess

returns on the market's excess returns, can be rendered insignificant by diversification, and, therefore, should not be rewarded by excess return. Systematic risk, measured by the slope coefficient, or beta, is, by contrast, inescapable and should earn a risk premium.

Drawdown In the words of Grant (2004),

If you stick with this business of trading and investment, know this: You are going to spend most of your time in an unhappy state known as drawdown.

Drawdown is peak-to-trough decline during a specific period, usually expressed as a percentage drop. It follows from the definition that, except when reaching a new high for the period, a portfolio is in a drawdown. This fact is illustrated in Figure 5.3: Except as reported by by Bernard Madoff, zero drawdown is usually achieved only on a small number of individual days. Most of the time, the drawdown is negative, its magnitude signifying the loss relative to the previous peak. (By definition, drawdown cannot be positive.)

This measure is most commonly applied to hedge funds, which, in an ideal world, if indeed properly and dynamically hedged, should always lock in gains. In practice, of course, continuous hedging is infinitely expensive, and even the best fund managers routinely experience declines from the peak. The objective of drawdown analysis is to assess the manager's ability to keep most of the gain in diverse investment environments, especially during periods of market disruptions.

Hedge fund risk managers would often establish maximum acceptable drawdowns for a range of trailing periods, for instance, one month, six months, one year, lifetime. If a trader's portfolio breaches one of the limits, the risk manager may scrutinize the positions and market environment and, depending on the circumstances, advise (or force) the trader to liquidate or immunize some or all risky assets.

When analyzing a portfolio's drawdown record, the analyst should consider not only the magnitude of the peak-to-trough decline but also the *recovery period*, namely, the time it takes the portfolio to recoup the loss since a previous high. In the bottom chart in Figure 5.3 recovery periods are the intervals between successive zeros.

Drawdown is an *extreme value* risk measure: It is the maximum (extreme) loss realized since the previous peak. This means that it probes the tails of the distribution of returns realized during the selected period. Unlike the mean, the variance, higher cumulants, and other *robust* risk measures, the drawdown is sample-dependent and does not converge to the




FIGURE 5.3 Daily value and drawdown for a hypothetical portfolio over a period of three years.

“true” value with increasing number of observations. On the contrary, as the observation period grows, the *maximum observed drawdown* would probe deeper into the tails and, in general, increase. Although challenging to treat analytically, drawdown is an important model-free empirical measure of a manager’s skill.

Typically, analysts calculate drawdowns based on daily portfolio marks. But they can also use lower frequency (e.g., weekly or monthly) when daily data are not available or suffer from significant serial correlations. We emphasize that drawdown is a function of both data frequency and the selected time period.

Return over Maximum Drawdown Defined as annualized period return divided by the worst drawdown experienced during this period, return over maximum drawdown (RoMaD) is a particular kind of risk-adjusted return, alongside Sharpe, Sortino, and Treynor ratios. RoMaD’s strength is that it scales the return with the actual *maximum* loss of portfolio value, not a measure of typical or probable loss or a degree of co-movement with the broader market. An attendant disadvantage is the difficulty of analytical treatment or forecasting.

Relative Risk Measures

A perfect hedge is rarely plausible and, as discussed in section 2.5, may not be desirable. A risk manager, therefore, needs tools for assessing the degree to which a hedging instrument mimics the behavior of the portfolio or asset being hedged. In other words, the risk manager needs to quantify *basis risk*. By contrast, a portfolio manager in search of a diversifier would be interested in an instrument that does not move in lockstep with the current investments and would need to quantify this property. Finally, passive and active fund managers typically measure their portfolios’ performance against a predefined benchmark, usually a market index or a combination of indices. Fund managers, their risk managers, and fund investors need a means of determining to what degree the fund deviates from its benchmark.

The statistics introduced below provide quantitative measures of two assets’ degree of co-movement and their relative risks. For convenience, we will often talk about a portfolio and its benchmark but the discussion applies to any two financial assets.

Tracking Error This is the standard deviation of the portfolio’s returns in excess of its benchmark, usually annualized. Let $\{r_i\}_{i=1}^N$ and $\{r_i^{(b)}\}_{i=1}^N$ be time series of returns on, respectively, a portfolio and its benchmark and

let $\Delta r_i \equiv r_i - r_i^{(b)}$. Then the tracking error τ is defined, by analogy with equation (5.1), by

$$\tau^2 = \mathcal{A} \times \frac{1}{N-1} \sum_{i=1}^N (\Delta r_i - \overline{\Delta r})^2, \quad \overline{\Delta r} = \frac{1}{N} \sum_{i=1}^N \Delta r_i, \quad (5.7)$$

$\mathcal{A}$ being the annualization factor. Note that tracking error is *symmetric*: Its magnitude does not change if one reverses the roles of the portfolio and the benchmark.

A good hedge would have a low tracking error. A passively managed index fund should have a tracking error much smaller than the benchmark index's volatility $\sigma^{(b)}$ (note that, as a consequence, the volatilities of the fund and the benchmark should be very close). An actively managed fund's τ should be smaller than $\sigma^{(b)}$ but large enough to allow sufficient deviation from the benchmark composition to deliver outperformance. At the same time, an asset with a low tracking error would not provide good diversification for a portfolio.

Volatility Ratio Deceptively simple, this measure is just the ratio of the portfolio's volatility σ to the volatility of its benchmark $\sigma^{(b)}$:

$$F = \frac{\sigma}{\sigma^{(b)}}. \quad (5.8)$$

Easy to calculate, it provides a quick test of the relative riskiness of two assets. As mentioned above, for both passively and actively managed funds, F should be close to unity. A large volatility disparity between a fund and an asset, however, does not by itself indicate a good diversifier: Assets with vastly different volatilities can be highly correlated, while a good diversifying instrument, weakly or negatively correlated to the portfolio, can be similarly volatile, that is, have $F \approx 1$.

Volatility ratio is particularly important as one of the two factors defining beta, as described below. The beta, or the slope regression coefficient, determines the size of the hedge. We will elaborate on this later in this section.

Unlike tracking error, volatility ratio is not symmetric: Interchanging the portfolio and the benchmark would replace a volatility ratio with its inverse.

Covariance and Correlation Coefficient Covariance is a measure of two assets' co-movement. Given time series $\{r_i\}_{i=1}^N$ and $\{r_i^{(b)}\}_{i=1}^N$ of returns, the sample covariance is

$$C = \mathcal{A} \frac{1}{N-1} \sum_{i=1}^N (r_i - \bar{r}) (r_i^{(b)} - \overline{r^{(b)}})$$

$$\text{with } \bar{r} = \frac{1}{N} \sum_{i=1}^N r_i, \quad \overline{r^{(b)}} = \frac{1}{N} \sum_{i=1}^N r_i^{(b)}. \quad (5.9)$$

If r and $r^{(b)}$ exactly match (i.e., $r_i = r_i^{(b)}$ for all i), then the covariance reduces to variance, equation 5.1. In particular, except in the degenerate case of $r_i = \bar{r}$ for all i , the covariance of identically moving assets is always positive. More generally, a positive covariance indicates predominantly similar moves and a negative one moves typically in opposite directions, although the degree to which the two assets' movements are synchronized is not readily apparent from the value of C alone. Similarly, while a small covariance magnitude signifies a lack of coordination between r_i and $r_i^{(b)}$, the precise meaning of "small" is unclear without further analysis.

These limitations are addressed by the *correlation coefficient*, defined as covariance, scaled with the two assets' standard deviations:

$$\rho = \frac{C}{\sigma \sigma^{(b)}}. \quad (5.10)$$

In contrast to covariance, the value of a correlation coefficient is confined between -1 (perfect anti-correlation) and $+1$ (perfect correlation), unity providing a magnitude scale. For example, a value of $\rho = 0.95$ instantly tells an analyst that the assets would mostly move in tandem, while $\rho = -0.5$ indicates that more often than not a positive r_i would coincide with a negative $r_i^{(b)}$. A correlation coefficient with a magnitude much smaller than unity is a sign that the two assets are uncorrelated.³

It is clear from their definitions that covariance and correlation coefficient are symmetric, that is, do not change if one swaps the asset and the benchmark.

Beta Beta relates returns on a portfolio or asset and returns on its benchmark. Mathematically, it is the slope coefficient in a linear regression: Given time series of returns, one finds α and β that minimize the sum of squared errors $\sum_i \epsilon_i^2$ in the linear model

³See Appendix B for a discussion of statistical significance of these statements.

$$r_i = \beta r_i^{(b)} + \alpha + \epsilon_i, \quad i = 1, \dots, N. \quad (5.11)$$

The idiosyncratic part of r , namely, α , is unrelated to the benchmark, while the systematic component of the return, $\beta r^{(b)}$, is driven by $r^{(b)}$ and can, therefore, be hedged.

It is straightforward to show that for a univariate regression the beta is a product of correlation coefficient and volatility ratio:

$$\beta = \rho F \equiv \rho \frac{\sigma}{\sigma^{(b)}}. \quad (5.12)$$

Writing the beta in this way highlights one of its important properties: Since F is not symmetric with respect to interchanging the asset and its benchmark, neither is the beta. This observation, along with the decomposition of equation 5.12, plays an important role in the practical method of hedge sizing, described in the next section.

5.3 OPTIMAL HEDGE SIZING

Suppose our portfolio has an exposure to a regional equity market, let us say, to be specific, Far East, proxied by a broad equity index, such as MSCI Far East, with an equivalent variation investment (EVI) value of $\mathcal{E}$. Suppose also we want to hedge this exposure. Unfortunately, the MSCI Far East Index is illiquid and shorting it is impractical. What can we do?

We may have determined that returns on the MSCI Far East have historically reasonably matched the returns on the NIKKEI 225, a liquid and easily tradable index (assume that we have converted both indices into U.S. dollars). But what should be the notional amount U of the short NIKKEI 225 position to provide an optimal hedge?

Our discussion of the beta at the end of the previous section might suggest that we should regress the hedge returns r_H (in our example returns on the NIKKEI 225) on the returns r of the reference index (MSCI Far East), that is, $r_H = \beta r + \alpha$, and use β for hedging. We want the change in the MSCI Far East EVI during a given period, namely, $r\mathcal{E}$, to equal the change in our hedge position, or $r_H U$, that is, we want $r\mathcal{E} = r_H U$. Substituting $r_H = \beta r$ (and disregarding the unhedgeable idiosyncratic portion of the return, α), we obtain

$$U = \beta^{-1} \mathcal{E}. \quad (5.13)$$

But we could have with equal justification regressed the reference index on the hedge, $r = \hat{\beta}r_H + \hat{\alpha}$, and, following similar logic, concluded that the hedge notional should be

$$\hat{U} = \hat{\beta}E. \quad (5.14)$$

Which answer, if either, is correct?

Let us have a closer look at the two hedge size ratios, $h \equiv U/E$ and $\hat{h} \equiv \hat{U}/E$. Using the beta decomposition of equation 5.12, we can write them as

$$h = \rho^{-1}F, \quad \hat{h} = \rho F \quad (5.15)$$

where ρ is the correlation coefficient between the reference and the hedging indices and $F \equiv \sigma/\sigma_H$ is these indices' volatility ratio. The only difference between h and $\hat{h}$ is in the way they involve ρ .

Let us digress and ask what properties an ideal hedge would possess. The answer is that it should be perfectly correlated with the reference asset, that is, the returns correlation coefficient should equal unity. The volatilities do not matter: If a perfectly correlated hedge is 10 times more volatile than the reference index, we would simply short one-tenth the EVI. (All of this, of course, assuming the usual requirements for linear regression are satisfied; see, e.g., Hamilton, 1994.)

Now note that when $\rho = 1$ the expressions for h and $\hat{h}$ coincide: They both reduce to the volatility ratio, the perfect hedge limit. In fact, the deviation of h/F from unity is a measure of the hedging instrument's suitability for the task. Let us assume that our hedge is not monstrously inadequate, that is, that $\rho = 1 - \delta$ with δ being a small positive number (say 0.1, or 0.2, or even 0.4). Then we can rewrite hedge size ratios, divided by F , in a form that explicitly displays the deviation from the perfect-hedge value of unity:

$$\frac{h}{F} \equiv \rho^{-1} = 1 + \rho^{-1}\delta, \quad \frac{\hat{h}}{F} \equiv \rho = 1 - \delta. \quad (5.16)$$

Observe that (assuming $\rho \neq 1$) $h > F$, $\hat{h} < F$, and, as a consequence, $h > \hat{h}$: regressing the hedge on the asset always yields a higher hedge notional.

Equation 5.16 offers a quantitative measure of a hedge's deviation from perfection: This deviation is of order $\delta = 1 - \rho$. It is not difficult to notice that, since $\rho^{-1} \equiv (1 - \delta)^{-1} \approx 1 + \delta$, the linear (in δ) terms in h and $\hat{h}$ are equal and opposite. This observation suggests that the quality of

hedging could be improved by using an average of h and $\hat{h}$ for sizing. Indeed,

$$h_{\text{opt}} \equiv \frac{1}{2}(h + \hat{h}) \equiv \frac{1}{2}F(\rho^{-1} + \rho) = F\left[1 + \frac{1}{2}\rho^{-1}\delta^2\right], \quad (5.17)$$

that is, the hedge imperfection is now of order δ^2 , much smaller than δ , characterizing both h and $\hat{h}$.

We can now summarize the steps needed to determine the optimal size of an imperfectly correlated hedge:

- Step 1. Calculate the slope coefficient β in the regression of the hedge on the hedged exposure.
- Step 2. Calculate the slope coefficient $\hat{\beta}$ in the regression of the hedged exposure on the hedge.
- Step 3. Short the notional amount

$$U_{\text{opt}} = \frac{1}{2}[\beta^{-1} + \hat{\beta}]\mathcal{E} \quad (5.18)$$

where $\mathcal{E}$ is the hedged exposure.

The error, resulting from less than perfect correlation between the hedging instrument and the exposure, would then be of order $(1 - \rho)^2$ where ρ is the correlation coefficient of returns on the hedge and returns on the hedged exposure. As ρ approaches unity, β^{-1} and $\hat{\beta}$ converge.

5.4 TAIL-RISK MEASURES

Tail-risk measures consider the likelihood and consequences of rare events. This means that, unlike cumulants, such as the mean, the variance, the skew, and the kurtosis, they reflect the properties of the *tails* of the distribution of returns rather than its bulge region, more specifically, the left, or downside, tail. Like cumulants (and unlike extreme value statistics, such as drawdown or RoMaD), tail measures are *robust* in the sense that they are a property of the returns' probability distribution and, as the sample size increases, their calculated value converges to a theoretical limit, completely determined by this distribution. Contrast this with drawdown: While dependent on the governing distribution, a drawdown is a property of the *sample*. Thus, the period maximum drawdown, as an example, would generally increase (it

can sometimes remain unchanged) as the period grows and the number of observed returns increases.

By their nature, tail-risk measures are supposed to address precisely the kind of rare events that inhabit risk managers' (and prudent investors') nightmares: "once-in-decades" multiple-sigma moves; "theoretically" impossible but brazenly real developments; counterintuitive confluences of events that go against "robust" historical relationships. The lamentable truth, however, is that, as typically used in financial industry, tail measures often simply restate the limited vision of risk contained in "bulge" risk statistics of section 5.2. Indeed, the most common practice is to calculate tail measures under the assumption of a normal distribution of logarithmic returns with the returns' mean μ and standard deviation σ either calculated directly from historical observations or deduced via a model. This approach adds nothing to the information already contained in a normal distribution's two defining moments and simply rephrases, perhaps more conveniently, the limited risk view implied by μ and σ . The empirically attested asymmetry of returns and presence of heavy tails are ignored.

There is, to be fair, a fundamental conceptual limitation that can explain, at least in part, why tail measures' risk-assessing potential is seldom fully utilized: Tail statistics endeavor to characterize events evidenced by few, if any, empirical observations. Consider, for example, a time series of five years of daily returns (some 1,300 observations). If one is interested in downside tail events that can occur in a single day with a 1% probability, then one has to draw conclusions from only about $0.01 \times 1,300 = 13$ empirical data points. The very number of observations in the 1% left tail can be determined only to within approximately $\sqrt{13} \approx 4$, or with a 30% accuracy (*cf.* equation (3.1)). This means in practice that an analyst cannot even tell whether he should consider nine or 17 lowest returns to represent tail events accounting for 1% of outcomes. Although in principle an analyst should be able to determine one-day 99%-confidence level value at risk (defined below) by simply noting the loss (or, as the case might occasionally be, profit) of the 13th worst data point, in practice the statistical error inherent in this determination would render the result virtually meaningless. An additional complication is that, since distant tails in a financial time series are hit infrequently (these are the "outliers," after all), there is a good chance that the market conditions, and, therefore, the probability distribution that governs the returns, change between hits.

This paucity of tail-event observations necessitates introducing assumptions about the shape of the distribution of returns. The standard procedure is to postulate the functional form of the probability distribution, fit its parameters to empirical data, and then calculate tail statistics analytically. The functional form used almost invariably in the financial industry is the

fast-decaying normal (Gaussian) shape, and the fitting consists in plugging in the realized or model-deduced mean and standard deviation of logarithmic returns. This approach systematically underestimates tail risk while creating a dangerous illusion of safety.

An analytical (i.e., distribution-derived, as opposed to time series-based) calculation of tail risk measures makes more sense under the assumption of a long-tailed Lévy distribution of returns, introduced in section 3.3. As discussed there, high frequency of tail events is immanent in Lévy distributions, and, consequently, would also be inherent in the derived tail measures. An additional advantage is that Lévy distributions, unlike the normal one, can be asymmetric and, therefore, able to capture skew risk.

Note also that whereas bulge risk statistics become meaningless (infinite) when Lévy distributions are assumed, value at risk, the most commonly used tail-risk measure, and median tail loss (both defined below) remain well defined and finite.

We describe and discuss below three commonly cited tail risk measures: *value at risk*, *expected tail loss*, and *median tail loss*. They can be calculated analytically, that is, from the presumed distribution of returns, or empirically, from historically observed performance. We urge the reader to use these measures with caution and be mindful of their intrinsic limitations. Dowd (2002) discusses advanced statistical tools devised to make the most of sparse data, available for calculating tail risk measures.

Value at Risk

Consider a portfolio whose value at the close of business on a particular day, day 0, is V_0 . Suppose we have an estimate of the probability distribution of possible gains or losses $\Delta V \equiv V_1 - V_0$ on the subsequent day, day 1. The estimate for this probability distribution $p(\Delta V)$ could have come from a model, a historical performance analysis, or been simply a guess. The portfolio manager might then want to know the maximum loss his portfolio can sustain on day 1 if the realized ΔV does not fall in $\lambda = 5\%$ of least favorable outcomes. The answer is called *value at risk* (VaR) *for one day time horizon at a $1 - \lambda = 95\%$ confidence level* (CL).

Note that value at risk is a function of two parameters: (1) *time horizon*, also known as *holding time*, and (2) *confidence level*. (Sometimes λ , not $1 - \lambda$, is called confidence level.) VaR for, say, one week or one month, for the same value of λ , would be different from (and, in general, greater than) VaR for one day. Likewise, a one-day VaR at CL = 99% would be greater than the VaR at CL = 95%.

The concept of VaR is graphically illustrated in Figure 5.4, which shows the probability distribution of one-day changes in the value of a hypothetical




FIGURE 5.4 Probability distribution function of one-day changes in the value of a hypothetical portfolio. Dashed vertical lines mark VaR at indicated confidence levels (CL).

portfolio. In order to determine value at risk ΔV_λ , corresponding to a particular confidence level $CL \equiv 1 - \lambda$, the risk manager needs to find a point $-\Delta V_\lambda$ on the horizontal axis such that the area, contained under the probability distribution function $p(\Delta V)$ to the left of $-\Delta V_\lambda$, equals λ . In other words, the probability that the loss during the stated time horizon (one day in this case) equals or exceeds ΔV_λ should equal λ , that is,

$$F(-\Delta V_\lambda) \equiv \int_{-\infty}^{-\Delta V_\lambda} d\Delta V p(\Delta V) = \lambda \quad (5.19)$$

where $F(\Delta V)$ is the cumulative distribution function (CDF). One can formally express the solution to this equation in terms of the CDF's inverse $F^{-1}(\cdot)$:

$$\Delta V_\lambda = -F^{-1}(\lambda). \quad (5.20)$$

This formula is graphically illustrated in Figure 5.5.




FIGURE 5.5 Cumulative distribution function of one-day changes in portfolio value, corresponding to the PDF in Figure 5.4. The arrows illustrate the steps in calculating value at risk ΔV_λ for a given confidence level $1 - \lambda$.

For many standard probability distributions, most notably the normal one, the CDF and its inverse are readily available in most standard statistics packages, and equation (5.20) can be used directly to calculate VaR for any λ . For more exotic distribution functions, distributions, derived from historical performance, or those, implied by Monte Carlo⁴ and other models, equation (5.19) can often be solved numerically. In either case, this equation expresses the basic concept of value at risk.

We emphasize that the probability distribution $p(\cdot)$ and, as a consequence, its cumulative distribution function $F(\cdot)$ and inverse CDF $F^{-1}(\cdot)$ are *time horizon-dependent*, and as time horizon changes, they change too, as does, via equation (5.20), the VaR.

Incidentally, the distribution of returns appearing in Figure 5.4 demonstrates the insidious dangers of tail risk. The most probable one-day return for the shown distribution happens to be on the order of positive \$1 million.

⁴For efficient VaR calculation using the Monte Carlo method see Glasserman, Heidelberger, and Shahabuddin (2001).

However, the distribution is negatively skewed and possesses a long left tail. As a result, there is a 1% probability of losing more than \$7.8 million, a 5% probability of a loss in excess of \$2.8 million, and a 10% probability of losing \$1.2 million or more. The likelihood of sustaining a loss of any magnitude is more than 20%. An analyst, blithely assuming normality and using the mean and standard deviation as inputs, would obtain much smaller VaRs and dangerously underestimate tail risk.

Note that while the typically reported values at risk are positive, in some cases, especially at lower confidence levels (e.g., 90% or less), a VaR can be negative. When this happens, the VaR's absolute magnitude should be interpreted as the minimum gain achievable at the stated confidence level. For example, if a portfolio has a VaR of $-\$1$ million at a 90% confidence level, then the chances of a gain of \$1 million or more are 90%. With a probability of 10% the portfolio would return less than \$1 million or suffer a loss.

Value at risk states the maximum possible loss at a given confidence level. It does *not* define the maximum loss that can possibly occur, nor does it say anything about the severity of tail loss should it occur. Suppose, as an example, that a portfolio's one-day VaR at 95% confidence level is \$1 million. This means that with 95% probability the loss on a given day will not exceed \$1 million. But this also means that there is a 5% chance of losing more than \$1 million in a single day. Moreover, if such a loss occurs, it can in general be unlimited. Apart from the obvious statement that by definition the loss will be greater than value at risk (\$1 million in our example), the likely magnitude of the loss will not be directly related to VaR.

We describe in the remainder of this section two risk measures, *expected tail loss* and *median tail loss*, frequently used to gauge the likely severity of a tail event.

Expected Tail Loss

Expected tail loss (ETL) is the mean loss magnitude given that a tail event occurs during a specified time horizon. The tail event is defined as a loss exceeding a preset threshold level ΔV_λ . Since there is a one-to-one correspondence between ΔV_λ and confidence level $1 - \lambda$ via equation (5.20), one can alternatively view ETL as a function of time horizon and confidence level (or, equivalently, λ). ETL is closely related to value at risk and, in fact, complements the concept of VaR by quantifying the likely loss size.

The main drawback of ETL is that, because of its dependence on the *shape* of the far tail of the distribution of returns, it is afflicted by the problem

of sparse empirical data even more than VaR. In addition, unlike VaR, it can become unhelpfully infinite if the left tail is governed by a Lévy distribution.

Given a probability distribution $p(\cdot)$ of returns (for a given time horizon), expected tail loss L_λ at confidence level $1 - \lambda$ equals minus the mean return, calculated with a renormalized truncated probability:

$$L_\lambda = - \int_{-\infty}^{-\Delta V_\lambda} d\Delta V \Delta V p_\lambda(\Delta V) \quad (5.21)$$

where the tail probability distribution $p_\lambda(\cdot)$ is defined as

$$p_\lambda(\Delta V) \equiv \begin{cases} \lambda^{-1} p(\Delta V), & \Delta V < -\Delta V_\lambda \\ 0, & \Delta V \geq -\Delta V_\lambda \end{cases} \quad (5.22)$$

with value at risk ΔV_λ related to λ via equation (5.19) or, equivalently, equation (5.20). In words, the shape of $p_\lambda(\cdot)$ is determined by the tail of $p(\cdot)$ to the left of $-\Delta V_\lambda$; the renormalization factor λ^{-1} simply ensures, via equation (5.19), that the new probability is normalized to unity. Figure 5.6




FIGURE 5.6 Tail probability distribution for $\lambda = 10\%$, corresponding to the distribution of returns in Figure 5.4. Dashed vertical lines mark expected tail loss and median tail loss at CL = 90% ($\lambda = 10\%$) or, equivalently, for a tail event loss threshold of $\Delta V_{10\%} = \$1.2$ million.

shows the tail distribution for $\lambda = 10\%$, corresponding to the probability of returns in Figure 5.4. A dashed vertical line indicates the one-day ETL at a 90% confidence level.

Median Tail Loss

Median tail loss (MTL) is the median size of the loss if a tail event occurs. That is, should a portfolio suffer a loss in excess of a preset threshold, the magnitude of this loss would equally likely fall below or above the MTL. As is ETL, the median tail loss is a function of time horizon and either confidence level or loss threshold.

A median tail loss is calculated from the same tail probability distribution introduced for evaluating the expected tail loss and illustrated in Figure 5.6. A dashed vertical line in this figure marks the one-day MTL at a 90% confidence level.

Being a median, the MTL is more robust than the ETL, a mean. This means that a calculated MTL value is less sensitive to outliers and statistical fluctuations. An additional convenience is that, unlike the ETL, the MTL remains finite and meaningful under a Lévy distribution assumption.

As is the case with the ETL and the VaR, the accuracy with which the median tail loss can be determined is limited by sparse statistics. Another limitation shared by all the three tail risk measures is that, while shedding some light on the distribution of extreme losses, none of them enlightens the analyst as to the maximum possible loss. Finally, a risk manager should not lose sight of the fact that, like most risk measures, tail statistics, whatever the calculation method, ultimately reflect observed past performance and, therefore, can fail when confronted with market conditions not experienced before.

CHAPTER 6

The Decision to Hedge

Stay! though the woods are quiet, and you've heard

Night creep along the hedges. Never go

Where tangled foliage shrouds the crying bird,

And the remote winds sigh, and waters flow!

Rupert Brooke, *Town and Country*

6.1 TO HEDGE OR NOT TO HEDGE?

In Chapter 2 we discussed many of the risks to which a portfolio can be exposed. We now turn our attention to the hedging of those risks and consider the elements in the decision about whether those risks should, and can, be hedged.

Until quite recently, when evaluating investment alternatives, finance professionals commonly subscribed to modern portfolio theory (MPT), which came in vogue in the 1960s. In recent years, the theory, which suggests that an efficient portfolio can be constructed that will, through diversification, maximize return for a given level of risk, has come under attack, with several of its assumptions being challenged, perhaps most notably those that market returns are normally distributed and correlations are constant. From a practical viewpoint, many proponents' faith in the theory was shaken when they suffered losses during the recent market turmoil, notwithstanding the vaunted benefits of diversification, and there is now a growing acceptance of alternative theories such as that of behavioral finance. One of the core propositions of MPT, however, that in rational markets investors seek increased returns for assuming incremental levels of risk, is still generally accepted to be true. The assumption is that while risk-averse investors can choose to eschew risk by investing in *relatively* riskless assets such as U.S. Treasury bonds and earn a "risk-free" return, more adventurous investors

will tend to invest in riskier assets, but will expect to be rewarded with commensurately higher returns. These topics can be explored in depth in any standard finance textbook.

Those who are exposed to market risk, be it as the result of holding a financial instrument or as a consequence of business activity that is exposed to commodity, currency, or other factors, would do well to heed the words of Benjamin Graham. In his classic text, *Security Analysis* (Graham and Dodd, 2002), he distinguished between *investing* and *speculating*:

An investment operation is one which, upon thorough analysis, promises safety of principal and an adequate return. Operations not meeting these requirements are speculative.

In another book, *The Intelligent Investor* (Graham, 1986), he admonished the reader:

If you wish to speculate, do so with your eyes open, knowing that you will probably lose money in the end; be sure to limit the amount at risk and to separate it completely from your investment program.

The purpose of a hedging strategy is to mitigate the risks to which an enterprise or portfolio is exposed, and its successful implementation can transform what might otherwise be a speculative venture into a judicious investment. In most cases, a hedging strategy—or, in short, a hedge—involves entering into a position or purchasing an instrument that has a risk profile that offsets, ideally directly, the exposure of the position to be hedged. Since the hedge itself has a risk exposure, financial derivatives and other instruments used for hedging purposes can also be employed, wittingly or unwittingly, for speculative purposes. An *appropriate* hedge is one that reduces the entire risk of the position, taking into consideration any increase in risks such as basis risk introduced due to an imperfect hedge, rather than one that seeks to maximize profit through interest rate, currency, commodity, or other speculation.

As discussed in Chapter 2, investors often measure risk in terms of variability of returns, using a measure such as standard deviation. We noted that many, although not all, risks can be mitigated using financial derivatives, that is, instruments or contracts whose value *derives* from some other, *underlying*, asset, index, or event. If a firm's hedging strategy proves effective, then risk—as measured by the variability of returns—will be reduced, resulting in turn in a reduction in the expected return that will be required by investors. Since the value of a firm can be determined by discounting its future expected cashflows at its required rate of return, a successful hedging

strategy that reduces that required rate of return should lead to an increase in firm value. However, the first rule of hedging, as we might define it here, is that there is no such thing as “a free lunch,” there being a cost associated with hedging so that, in an efficient market, any increase in firm value that derives from a lower cost of capital resulting from the successful implementation of an effective hedging strategy would be exactly offset by the cost of that strategy.

Since a consequence of theories such as MPT is that, in a perfect world, companies are priced according to their risk and return, there should be no advantage gained by a company choosing to hedge its exposures. In practice, however, the capital markets are not perfect, company management is better informed than shareholders about the risks a company faces—there is an *information asymmetry*—and hedging can be inefficient for individuals because of transaction costs. Consequently, it is the management of a company who are generally best qualified to evaluate the risks faced by the company, to decide if those risks can be effectively and economically hedged and, if so, to implement a hedging strategy.

Reasons to Hedge

For many enterprises, the primary goal of hedging using financial derivatives is to reduce the volatility of expected future cashflows. As discussed above, if markets are truly efficient, then no net benefit accrues to a firm that chooses to hedge. However, in practice there are a number of reasons why hedging can be a worthwhile exercise for a firm.

First, if a firm's profits are subject to a progressive tax scale, with a higher marginal tax rate being applied to higher levels of profits, then a company with more volatile earnings will, on average, pay more taxes than a firm that has the same earnings but with less volatility. If earnings volatility can be smoothed through hedging, the tax benefits may outweigh the cost of hedging.

Second, the management of a firm is likely to have more information than an outsider, especially an individual investor, about the risk profile of a firm, because of an information asymmetry, and so should be better placed to design and implement a hedge, especially a delta-hedging strategy, whose successful execution depends on a timely and accurate knowledge of positions and risk sensitivities. In addition, transaction costs for a firm will be lower than for an individual investor who seeks to hedge only a portion of a firm's total risk exposure.

Third, a firm's lenders may require, as a condition for lending money, that the firm hedge. This prerequisite may be stated as a *covenant* to a loan agreement, which dictates the terms under which a bank is willing to lend

money and the terms and conditions with which a firm that wishes to borrow must comply.

Fourth, by implementing an effective hedging strategy, a firm will reduce short-term cashflow volatility and lower the expected *cost of financial distress*. If a firm were to enter financial distress, or to go bankrupt, it would incur high incremental direct costs, such as legal and accounting fees, and indirect costs due to the actions of other stakeholders such as suppliers, customers, and lenders. The latter, for example, would be likely to significantly increase the price of lines of credit, or simply to terminate uncommitted credit lines. A hedging strategy that reduces the likelihood of company failure will, therefore, increase, or avoid decreasing, enterprise value. This can be a particularly relevant consideration in developing nations, where debt repayments are commonly linked to revenue from the sale of commodities, and hence directly linked to commodity prices. Furthermore, by reducing the likelihood of bankruptcy, a firm's creditworthiness will increase, which can lead to a corresponding decrease in the cost of borrowing.

Fifth, there may be a fiduciary responsibility to hedge, particularly in institutions such as pension funds. In the United States, the Employment Retirement Income Security Act (ERISA) incorporates a principle known as the "Prudent Expert Rule," which requires that a fiduciary exercise the same skill, prudence, and diligence when managing a portfolio as would a prudent professional in the conduct of a similar enterprise. In 1992, shareholders of an agricultural cooperative that suffered over \$400,000 in losses from grain sales successfully sued the directors of the co-op for their failure to adequately hedge financial risk. Although the directors had been aware of the importance of hedging, they failed to ensure that management implemented an effective hedging strategy. In its decision in the case of *Brane v. Roth*,¹ the Indiana Court of Appeals agreed with the shareholders, reasoning that since the bulk of the cooperative's revenues derived from grain sales, a failure to hedge was a breach of the directors' duty of care to the shareholders. For a discussion of this and related issues, the reader is referred to Chew (1996).

Sixth, a manager may choose to hedge for short-term, tactical reasons. He may have a fixed income portfolio, for example, that is intended to finance a future construction project. While the interest rate exposure of the portfolio may be one with which the manager is comfortable in the long run, he may deem it prudent to reduce the portfolio duration while there is some short-term uncertainty in the market, such as the outcome of a forthcoming Fed meeting that may change interest rate expectations. Rather than sell bonds and move to cash, with the intention of moving back

¹590 N.E. 2d 587 (Ind. Ct. App. 1992).

into bonds once the Fed decision has been announced, the manager could reduce interest rate risk while minimizing transaction costs by employing a hedging strategy such as entering into, for example, a short interest rate futures position.

These reasons suggest why, within the context of diversifiable risk, an enterprise may be justified in hedging. In practical terms, however, a portfolio manager in a large institution may have no other choice than to hedge, at the behest of a risk manager or because of a firm-wide policy, and will not care what the underlying rationale may be; rather, his energies will be directed toward structuring an effective hedge that balances the firm's *willingness* and *ability* to take risk. An individual investor, meanwhile, may be motivated to hedge by no other rationale than the desire to avoid outright speculation, and so may invest in, for example, a stock market index fund while simultaneously purchasing disaster insurance in the form of out of the money put options, or as an alternative may choose a *principal protected note*, designed to offer a combination of upside potential with safety of principal.

Comparative Advantage

Some firms will use the principle of comparative advantage in risk taking to decide which risks to retain and which to lay off. Embodied in this principle is the concept that a firm's management has a distinctive competence or an informational advantage over competitors in managing the risk of certain activities. The firm is better off retaining these risks and managing them internally, rather than seeking to hedge them with an external counterparty who might require a significant premium to assume the risk. Other risks, however, that the firm has no specialized competence in managing, such as interest rate risk and exchange rate risk, can generally be hedged cost effectively, leaving the company to focus its resources on managing its core activities. An airline company, for example, is in the business of transportation, not in the business of attempting to profit from movements in the price of oil. By hedging its exposure to oil prices, the airline company can stabilize its income and, in particular, reduce the volatility of the ticket prices it charges its customers.

One apparent disadvantage of hedging, whether intended to minimize the variance of expected returns or simply to provide "disaster insurance," is that the cost of the strategy will reduce expected returns in normal markets. A firm that hedges will tend to underperform in benign markets relative to those competitors who choose not to hedge and may lose customers and investors as a result. Just as fire insurance could be considered to have been a waste of money if your house does not burn down, similarly an airline that

“loses” \$50 million in the futures market as a result of hedging its fuel costs might be compared unfavorably, especially by naïve commentators in the press, to its competitors who have a larger risk appetite or who are financially less sophisticated. The danger for such a firm is that, in an environment where customers, investors, and competitors suffer from *disaster myopia*, it might not survive long enough to demonstrate the wisdom and superiority of its conservative strategy in adverse market conditions. Moral hazard can also play a role in the decision of how much risk to bear, particularly for financial institutions. While regulators might require that all banks maintain a prudent level of reserves to protect against adverse market conditions that might reasonably be expected to occur, they might not require reserves for the effects of crashes or crises that are expected to occur not more often than once every 80 years, for example. A bank that chooses to hold reserves to protect against such unlikely events will be uncompetitive in a marketplace where competitors implicitly believe that they are “too big to fail,” and who neglect to reserve against unanticipated losses in the belief that there is no need to.

Justification for Hedging

The rationale for hedging will vary from manager to manager. For some, it is a fundamental element of conducting business. If a firm makes use of a commodity such as copper as a raw material for its cookware products, for example, then its profitability will be critically dependent on the cost of the commodity. If the price of copper were to rise relative to iron or aluminum, the firm’s profit margin would be compressed if customers were likely to respond to increased prices for the firm’s copper products by switching to other types of cookware. However, if there were no alternative to a firm’s products, and competitors did not hedge, then a firm might be better off retaining the risk, since it could pass on the impact of increased raw material costs to its customers.

In many cases choosing to remain unhedged is an exercise in speculation. A manager who chooses not to hedge foreign exchange risk, for example, is speculating on rates remaining the same or moving in a favorable direction, with his decision perhaps being based on a forecast derived from the prevailing level of futures prices. Research suggests, however, that forward exchange rates are not a good predictor of future rates, as discussed in Solnik and McLeavey (2003). The reader is reminded that, even if markets prove to be mean-reverting, it may take some time for that phenomenon to manifest itself—by which time a firm may have become insolvent. An airline that chooses not to hedge the cost of its aviation fuel supply, for instance, may

come to rue this decision, if indeed the airline survives sufficiently long to have the luxury of regret. And to be perfectly clear, a manager who does not make a decision to hedge is implicitly making a decision not to hedge.

Markets for Risk

Some readers, recognizing that risk, like energy, cannot be destroyed, merely transferred—a principle that we may refer to as the law of conservation of risk—may ask how a market for risk can exist, and who, for example, would be willing to take on the risk that an airline is seeking to offload? In general, there are two markets for risk. First, for every player seeking to protect against one risk there are usually—but not always—other players in the same market, with differing needs and risk appetites, who are willing to assume that risk—for a price. In the case of the oil market, for example, market participants such as crude oil producers and oil refineries often seek to guarantee a minimum price floor for their products, while consumers such as airlines, power utilities and shipping companies usually seek to cap the price of the same products. The oil producer may be happy to receive, for example, a guaranteed \$70 per barrel for oil at the cost of foregoing even higher prices, while a road transport company in the United States might be happy to pay a fixed cost of \$2.70 a gallon for diesel, for example, content to lose the chance to benefit from a price decrease in return for being protected against the potential for increased prices. Similarly, oil refineries can make use of a *crack spread* to hedge against adverse relative price movements of crude and refined oil. The function of brokers and exchanges is to bring counterparties with opposing needs together. Banks, meanwhile, may find natural counterparties among their own customers. For example, a bank that has extended fixed-rate auto loans to its customers may be concerned about the interest rate risk that it faces when financing such loans using short-term, floating-rate debt. As a result, it may seek to enter into an interest-rate swap with a corporate client concerned about the interest-rate risk of its fixed-rate debt obligations. If the corporate client believed that interest rates were likely to decrease, it would be happy to enter into such a swap, undertaking to pay floating and receive fixed, while the bank would be happy to receive floating and pay fixed.

In addition to the natural players, however, there is also the speculator, who seeks to profit from supply and demand imbalances and market inefficiencies. It is not uncommon for investment firms and individuals to buy and sell crude oil, for example, with some speculators even owning their own storage and transport facilities. Speculators play a vital role in the markets by providing liquidity. While they often profit from their activities,

especially in times of high price volatility (when, ironically or not, risky option positions can increase dramatically in value), speculators sometimes come a cropper, with two notable casualties being the Hunt Brothers, who attempted to corner the market in silver in the late 1970s and were undone in 1980 by changes in exchange rules and margin requirements, and Amaranth Advisors, who in 2006 lost \$6 billion in a failed natural gas futures strategy.

6.2 THE HEDGING PROCESS

Once the decision to hedge has been made, at least tentatively, the design and implementation of a hedging strategy is a multi-step process, as illustrated in Figure 6.1, requiring the identification of both the risks to be hedged and an appropriate strategy that can be employed to hedge them. In most cases, it will be necessary to receive the approval of other constituencies in a firm, including the risk management and accounting groups, prior to implementing a hedging strategy, and coordination with the back office and middle office functions will be needed to ensure that any required accounts are opened, and that the appropriate systems and feeds are available to report on the new hedge.




FIGURE 6.1 Steps in the hedging process.

Identifying Risk Exposures

The first step in the process is for the manager to identify and quantify the risks to which the position or portfolio is exposed, as discussed in Chapters 2, 3, and 5. A prerequisite to this risk quantification is the ability to value the position using a robust model that makes use of prices and other market information delivered by a reputable financial data provider. Unless a position valuation can be produced in a consistent framework, and sensitivities to market factors determined, the construction of an effective hedging strategy can be no more than the product of guesswork.

Once a reasonable valuation for the portfolio can be produced, the decision to hedge the risks under consideration will often be predicated on an assessment of the magnitude of the economic loss incurred if markets were to move unfavorably. An example of such a loss estimate is *value at risk* (VaR), defined in section 5.4 as the maximum loss that should be expected, over a defined period of time, at a given level of confidence, under normal market conditions.²

To determine the appropriate size of a hedge position, the risk exposure must be quantified, and the characteristics of the assets or liabilities to be hedged must be identified. This is particularly true for commodities, for which the grade, position, size and, in some cases, delivery port must be specified. The liquidity of the underlying security or commodity should be determined, as should the shape of the appropriate forward curve, when applicable, both of which will help determine the appropriate hedging strategy. Next, the sensitivity of the portfolio valuation to changes in market factors must be calculated. These sensitivities, which, depending on the type of assets in the portfolio, may include delta, gamma, vega, DV01, and various duration measures, are described in Chapter 4. Factor sensitivities such as duration or delta can be easy to calculate for some assets, such as plain vanilla bonds or equity options. For other assets, however, such as complex asset-backed securities backed by mutual fund fees, it may be more practical to make use of proxies. This topic is explored in section 7.4.

The potential investment universe contains a myriad of instruments, with risk profiles ranging from what might be termed conservative to the wildly speculative, but all of which expose the investor to some degree of risk. For example, as discussed previously, even what is generally considered a “risk-free” investment in U.S. Treasury bonds in practice expose the investor

²It should be noted once again that VaR does not define the maximum loss that can possibly occur. At a 99% confidence level, for example, the probability that a loss will exceed the VaR is at least 1%, and quite possibly much higher if one allows for model uncertainties.

to a multitude of risks that include sovereign risk, inflation risk, reinvestment risk, and operational risk. And, as discussed in Chapter 2, investors and stakeholders in practically every business enterprise, from the simple farmer to a sophisticated conglomerate, are exposed to varying degrees of risk. This broad array of risks presents varying degrees of difficulty to those who seek to mitigate their exposures.

At one end of the spectrum are those exposures that can be easily hedged. For example, a U.S. exporter who is due to be paid in pounds sterling can enter into a foreign exchange contract to protect against the risk of the U.S. dollar appreciating against sterling. He can also protect against the credit risk of the obligor by, for example, requiring the debtor to provide a *letter of credit* from a highly rated financial institution that in effect guarantees the obligation. Similarly, a bond investor may choose to protect against the risk of default by a bond issuer by purchasing a credit default swap on that issue, or “name.”

At the other end of the spectrum are those risks that can be much more difficult to quantify and to hedge. Consider, for example, the risks faced by an insurance company that has sold an annuity. In return for a lump sum paid up front by the annuitant, the insurance company is obligated to pay a sum, generally a fixed amount, on a periodic basis to a recipient, usually the purchaser. In some cases, the payments are required to be paid for a finite term, but in many cases the obligation exists until the death of the annuitant. While death may be an inevitable end for us all, the timing of that unhappy event is almost impossible to predict with certainty for a given individual. The insurance company, therefore, is exposed to *longevity risk*, which refers to the risk that the annuitant will live longer than the insurance company had calculated when determining the price to charge for the annuity. Pension plans face a related risk, namely that the required payout on plan assets will be higher than modeled, a consequence of pension beneficiaries living longer than had been projected when the required size of pension contributions was being originally determined.

Insurance companies face another death-related risk, referred to as *mortality risk*, which relates to the risk that an insured person may die sooner than expected. An insurance company can expect to suffer a loss on a life insurance policy if the policyholder dies sooner than was projected using mortality tables. However, unlike a pension plan, an insurance company that provides both life insurance and annuity products may be able to benefit from a *natural hedge*—a position that by its very nature mitigates some or all of the risk of another position—if, on average, the pools of life insurance purchasers and annuity purchasers share similar demographic characteristics, such as lifestyle, geography, and socioeconomic status, and experience the same rates of mortality. In this case, the insurance company may assume

that changes in life expectancy will not be restricted to individuals but will be generally reflected in the population at large and that the mortality and longevity experience of the purchasers of both life insurance and annuity products will reflect that of the general population; it may decide, therefore, not to hedge these particular risks.

When analyzing risks, the manager must determine which, if any, positions provide a natural hedge to other positions, which risks ought to be hedged, and which risks cannot be hedged effectively and economically. We cannot overemphasize how critically important it is, when evaluating risk exposures, to scrutinize all of the assumptions used in the valuation model and in determining risk sensitivities. For example, any assumption of correlation between assets should be examined for reasonableness. The historical landscape is littered with the wreckage of firms that neglected to pay sufficient heed to the possibility that correlations between market factors would change dramatically over time. The use of *stress-testing*, discussed in greater detail in section 4.3, examining how the valuation of a position would change under market conditions of extreme stress, such as the equity market declines of October 1987 and August 1998 and the liquidity crisis of 2007–2010, can prove exceedingly insightful. Any hedging strategy that does not take extreme market conditions into consideration is fraught with danger. The brief discussion of portfolio insurance later in this chapter highlights an example of a hedging strategy that failed spectacularly in a time of market stress.

Risk Limits

The calculation of the VaR and other risk measures results in information on the magnitude of the risks to which the portfolio is exposed and estimates of the potential losses expected to be incurred for given changes in market conditions. A manager will generally have constraints on the amount of risk exposure that he is allowed to assume, and the second step in the hedging process is to specify these limits, which can be in the form of *stop-loss limits*, *notional limits*, or *exposure limits* (such as duration, DV01, or delta). The limits are usually established by a *risk manager*, who might be responsible for overseeing and managing the risk exposures of a number of trading desks, departments, or perhaps an entire business. The choice of limits will depend on the type of risk that the desk or business faces, the firm's tolerance for risk and profit and loss (P/L) volatility, and its level of expertise in managing risk, and regulatory constraints. Over time, as risk exposures change and as risk limits are increased or decreased, a manager considering the risk sensitivities and the VaR of his portfolio may decide to

close out or reduce some positions, to maintain other positions unhedged, and to hedge, in whole or in part, other positions.

Evaluating Alternatives

If the manager decides to structure a hedge, the third step in the process is to identify one or more potential hedging solutions and the counterparties who can provide the hedge. For relatively simple exposures, such as the foreign exchange exposure discussed above, a hedging solution can be created by establishing an offsetting position using a forward contract or a futures contract. For other exposures, the optimal hedge may involve swaps, puts or calls, zero-cost collars, or a combination of derivatives.

An appropriate hedge ideally will have the same underlying currency as the foreign currency to be hedged, since the price changes in the two positions should then offset. For liquidity or other reasons, however, it may not be practical or desirable to hedge the exposure directly, so another highly correlated currency may be used as a *proxy*. This is often the case when the currency to be hedged is that of a small country whose economy is highly dependent on that of a larger neighbor, or when the currency is *pegged* to that of one or more other countries. A company with exposure to the Swazi lilangeni, for example, may choose to hedge it using the South African rand, to which the Swazi currency is pegged. Alternatively, a company that has an obligation to pay in one currency and receive another, highly correlated, currency may choose to *net* the two exposures. In both of these cases, however, the companies are exposed to basis risk due to the nature of the *cross hedge*, which could result in an economic loss if the currency peg were to break or if the correlations between currencies were to change.

Basis risk can arise from a variety of sources. One cause can be a mismatch in assets, such as hedging a long position in a small number of stocks with a short position in a broad-market index future. Similarly, hedging one grade of oil with another introduces basis risk because the underlying commodities are not identical. Basis risk can also arise from a mismatch in the characteristics of long and short positions, such as hedging oil produced in Canada with a short futures contract that requires delivery of the same grade of oil but in a different location, such as Oklahoma. In the latter case, the obligation to deliver physical oil can be avoided by closing the position out early, for cash.

Many times, basis risk is considered as being acceptable in view of the cost savings compared with an alternative hedge. For example, hedging a small number of stocks with a short index futures position requires no initial cash outlay, unlike the alternative of purchasing put options on the individual stocks. An oil producer who wishes to hedge 10 years of production

of Dubai crude, for example, may prefer not to hedge with Dubai crude futures, which have poor liquidity, but may choose to hedge with a short futures position in a more liquid grade of oil such as Brent. The producer could subsequently enter into an additional contract such as a Brent/Dubai EFS (Exchange of Futures for Swaps) to mitigate the basis risk between the two grades of oil. An advantage of such a hedging solution is that Brent futures exist with a maturity of 10 years, whereas Dubai futures have a maximum tenor of three years.

Aside from basis risk, and the implications of changing correlations, any other assumptions underlying the efficiency of the proposed hedge should be carefully examined. In particular, if a perceived inefficiency in the marketplace makes one solution cheaper than an alternative, there may very well be inherent risks in the cheaper strategy that are not immediately apparent. For example, during the 2000s credit spreads on corporate bonds tightened as a result, in part, of a belief by many investors that they could safely hedge the credit risk of issuers through the use of credit default swaps. At the same time, an issuer of asset backed paper who could, for example, issue an “A-rated” bond at a spread of 100 basis points over swaps, was instead likely to exploit an inefficiency in market prices by choosing the alternative strategy of paying a “AAA”-rated monoline insurer a premium of 25 basis points to insure the bond, thereby raising its credit rating to a “AAA” level, for which investors would demand a credit spread of only 20 basis points, resulting in a saving of 55 basis points of coupon expense. What the issuer, in effect, was doing was exchanging the credit risk of an “A-rated” asset-backed security (ABS) structure, for which they would have had to compensate investors 100 basis points,³ for the credit risk of a monoline insurance company. What many investors failed to realize or pay sufficient heed to was the mispricing of risk on the part of monoline insurance companies, who may have understood the risk of each individual bond they were insuring—and, admittedly, quite likely at a level of understanding higher than that of individual investors—but who failed abjectly to appreciate the significance of the enormous concentration of correlated credit risk in their portfolios. Instead of purchasing a diversified ABS investment, investors in bonds wrapped by the monolines were buying credit insurance from a massively concentrated credit portfolio, resulting in their replacing one risk exposure, for which in the past they might very well have been adequately compensated, for one for which they did not receive sufficient compensation.

A jaundiced observer might also well ask why investors flocked, in the mid 2000s, to invest in collateralized debt obligation (CDO) structures,

³In practice, this spread also included a liquidity premium.

whose *raison d'être* seemed often to be merely to provide a means for investors to circumvent the restrictions in their investment guidelines that prevented their investing in lower-quality assets. After all, the same diversification benefits could be achieved by managers of all but the smallest investment portfolios by investing directly in the underlying bonds, without the necessity of paying the CDO management fee. One could argue, perhaps, although not very convincingly, that the CDO manager performed a necessary due diligence function that justified his fee. Not very convincingly, we say, because the performance of many of these CDO structures suggests that the investor could hardly have done a worse job selecting individual bonds himself.

It is of vital importance to step back, consider all the elements in the position and the proposed hedge, to systematically question each assumption used, and to consider the impact of changing market conditions, and the potential failure of counterparties, while bearing in mind what we may deem the second rule of hedging, “whatever can go wrong, will.” In 1998, for example, the hedge fund Long Term Capital Management (LTCM) had massive, leveraged, positions of \$125 billion (on \$4.7 billion of equity capital) focused predominantly on relative value, credit spread, and volatility strategies. These strategies were predicated on the assumption that market factors, such as bond yields and credit spreads, for example, that had deviated from their long-term averages were likely to revert to those long-term means. Unfortunately for LTCM, however, the Russian debt default in August 1998, followed by Brazil’s currency devaluation a short time later, caused a “flight to quality” that resulted in increased credit spreads and equity volatility. While LTCM had believed that its varied investment strategies provided a diversification benefit, justifying its use of high leverage, in fact those strategies were all based on the same assumption that risk premiums, and volatility, would decline. LTCM’s high leverage meant that even modest market losses could erode its equity position very quickly. In August 1998, it lost 44% of its capital, requiring the unwinding of positions, forced liquidations, and a vicious cycle of price declines. Ultimately, LTCM was a victim of a failure to thoroughly analyze the risks in its positions and was brought down by a combination of market risk and liquidity risk. History may demonstrate that the assumptions underlying the firm’s strategies were ultimately sound, but, as we pointed out in section 6.1, a firm’s reserves can disappear faster than the markets’ return to “normalcy.”

Counterparty Exposure

An important question to ask when hedging is, “Who is the counterparty?” The nature of the hedging instrument can make this a critical question. For

Forwards

- Customized
- Internal mark-to-market
- Traded over the counter (OTC)
- No exchange of cash until maturity
- Potentially significant counterparty risk
- Illiquid

Futures

- Standardized
- Official settlement
- Exchange-traded
- Daily margining
- Exposure is to clearing house
- Liquid, although limited selection

FIGURE 6.2 Comparison of futures and forwards.

example, if the instrument is a futures contract or an exchange-traded option, which have standard terms and conditions, the counterparty will be an exchange. Futures and options contracts are *margin*ed, requiring the deposit of both an initial performance bond and, subsequently, additional funds if there is an increase in the obligation owed by the holder of the contract. As a result of this daily settlement, counterparty credit risk, which is to the exchange, will be limited and, due to their standardized nature, there will be a liquid market for the instruments. In stark contrast, as shown in Figure 6.2, forward contracts, as well as other over-the-counter (OTC) derivatives, which are customized instruments, take the form of contracts between the purchaser and the seller directly, often with no exchange of cashflows between the parties until the expiration or maturity of the contract. Consequently, if the hedge instrument subsequently becomes in the money, or has a positive mark to market, there can be substantial counterparty credit exposure. The creditworthiness of the hedge provider, usually evidenced by its credit rating, becomes of paramount importance in the hedging decision. This is especially true in extreme market conditions, when a counterparty may fail precisely when a hedge is needed most. A company may have internal *counterparty limits* on the amount of exposure it is prepared to take against a given firm; if that limit has already been reached then a manager may be precluded from entering into a new hedge contract with the proposed firm, even if it offers a better price.⁴

To reduce counterparty risk, a firm may spread its business among several counterparties, although in reality the benefits afforded by diversification can prove ephemeral. For example, while the risk of default by a single “AAA”-rated counterparty over a five-year period might be calculated

⁴Counterparty limits are sometimes set at two levels; when exposure reaches the first level, no additional contracts with the counterparty can be initiated. When the second level is reached, existing contracts must be reduced in size or terminated or the counterparty must post additional collateral.

to be, say, 0.05%, there is a high degree of correlation between firms in the same industry, such as banks, so that the failure of one firm might very well be accompanied by the default of a peer. As a result, the probability of multiple defaults during times of stress is likely to be higher than expected. Furthermore, for a given firm the actual default risk may be much higher than suggested by its credit rating, since the frequency of crises is unknown in advance.

Most OTC contracts are based on templates provided by the International Swaps and Derivatives Association (ISDA). While trading on the OTC market, which consists of customized trades that are transacted by telephone or computer interactions between participants, can expose participants to greater liquidity and credit risk than the exchange-traded market, it nevertheless plays a vital role, allowing practically any exposure to be hedged, for a price. The volume of business conducted on the OTC market dwarfs that done on the exchange-traded market. In June 2009, for example, the OTC market was over \$600 trillion of notional principal outstanding, with interest rate swap contracts representing over 70% of the volume, whereas the exchange-traded market accounted for less than \$100 trillion.

Choosing the Optimal Strategy

The fourth step to be performed by a manager is to conduct a cost/benefit analysis on the proposed hedge alternatives. For example, due to their *non-linear nature*, options can provide protection against loss, while retaining the potential for profit, but at a cost, that of the option premium. *Linear* instruments, such as swaps, however, may provide protection with no initial outlay but at the cost of foregoing potential future profits. If there are a number of alternatives available to the manager, the advantages and disadvantages of each proposed strategy should be compared to allow an optimal solution to be chosen from among the available alternatives. The manager, operating within the constraints placed upon him by his risk manager and his risk/return objectives, must decide if the proposed hedge is acceptable and, if so, what the *exposure hedge ratio*, that is, the proportion of the exposure that he wishes to hedge, is to be.⁵ For example, a manager may decide to use a linear instrument such as a swap, but with a reduced notional, that

⁵We use the term *exposure hedge ratio* here to avoid confusion with the term *hedge ratio*, which can have the same meaning but is more generally used to refer to the ratio of the change in price of the hedging instrument to the change in price of the underlying. This topic is explored in Section 7.3.

is, less than the total exposure, in order to retain some upside potential. Of course, retaining the potential for some upside in this fashion also retains the potential for some loss.

The reader is cautioned that derivatives pricing is very often opaque to the purchaser. When deciding upon a hedge strategy, particularly one that involves the use of customized, OTC derivatives, a purchaser should seek quotations from a number of derivatives desks. This is true even when a manager is normally precluded, by firm policy, from trading with a desk outside his own institution. On occasion, internal quotes can be two or three times as expensive as an equivalent solution offered by a competitor.

Before entering into a hedge, a manager should be confident that the strategy will pay off in the right circumstances to prove effective. For example, a *knockout* put option will be cheaper than a plain vanilla put option, because in addition to the expiration date it will also have an *expiration price*, and therefore will not provide the same level of protection. Consider an investor who owns a stock with a price of \$100 and seeks to purchase protection against a price decline. If he is presented with two alternatives, a plain vanilla put option with a strike of \$100, or a down-and-out put with a barrier of \$85, he may very well choose the latter because of its cheaper price. However, the danger he faces is that if the stock price falls below \$85, the option will have hit its expiration price and so will have ceased to exist, with the investor having no protection precisely when it is most needed.

A manager should be aware of two other important hedge characteristics, *exercise style* and *path dependence*. The two most commonly used option exercise styles are *American*, which gives the owner a right to exercise on any business day before the expiration date, and *European*, exercisable only on the last business day before expiration. An intermediate style, *Bermudan*, named in reference to the island chain's geographic location between Europe and America, can be exercised only on a predetermined sequence of dates between purchase and expiration. Additional optionality costs money and, consequently, OTC options, which are custom-designed for the hedger's specific needs, are typically European. Standardized single-name exchange-traded options, intended for a broad range of investors with diverse objectives, are usually American.

The value of *path-dependent* instruments, as the name implies, depends not only on the instantaneous readings of the price of the underlying instrument, its (perceived) volatility, interest rates, and other market factors, but also on all the values the underlying price has assumed since inception. One path-dependent instrument, commonly used for hedging, is the Asian option. Its payout depends on the *average* value of the reference quantity (stock or

commodity price, index level, currency exchange rate, interest rate) over the life of the option. A typical user would be a business with foreign exchange exposures that wants to soften the possible negative impact of currency fluctuations. The business, however, may be concerned with average income volatility over the whole fiscal year rather than with monthly results. A single Asian option can hedge against the average exposure over the year while being much cheaper than a sequence of 12 one-month contracts and, requiring only a single decision, would incur less overhead. It would even be cheaper than a one-year European option on the spot exchange rate since the Asian option's smoothed underlying would be less volatile.

Other commonly used path-dependent instruments are the *look-back* option, whose payout depends on the maximum or minimum underlying price achieved during the option's lifetime, and the *shout* option, where the holder can "shout" to the writer, typically once, during the life of the option with the payout at expiration being the greater of the usual European option value and the intrinsic value at the time of the shout.

Look-backs are often written on commodities. A *floating strike* look-back call allows the holder effectively to buy the underlying at the lowest price $p_{\min}$, achieved during the life of the option, by, in essence, setting the strike price retroactively at maturity to $p_{\min}$. Similarly, a floating strike look-back put is a way for the holder to sell the underlying at the highest price $p_{\max}$ achieved while the option was in existence, by setting the strike retroactively to $p_{\max}$. (An alternative kind of look-back option, *fixed strike*, sets the strike price at inception but is exercised at the maximum or minimum, depending on the option type, level the underlying price has achieved during the life of the option.) Given the strong guarantee provided by floating strike look-back options, it is not surprising that they are very expensive. Shout options offer some of the features of look-backs at a considerably lower cost.

It is also important to understand how the option is priced, both to be able to estimate the dealer profit and to understand the P/L implications of market moves. Equally important is understanding how the position can be reversed, or closed out. For example, in *one-way markets*, customers attempting to close out a position may find that a dealer is ready only to sell, and so a customer seeking to terminate a position may discover it is necessary to negotiate a price—often punitive—with the dealer, one that may not reflect the true economic value of the instrument.

If the hedge involves the use of margined instruments such as futures, it is important to ensure that the firm has sufficient liquidity, which may be funded or may take the form of lines of credit, to be able to meet margin calls. In the next section we discuss the Metallgesellschaft case, where a lack of liquidity led to massive losses from a futures hedging strategy.

Approvals

Having designed a hedge and identified a suitable counterparty, it will be necessary to ensure that all required internal approvals have been obtained from the various authorities in the firm. If the hedge is constructed using instruments that have not previously been employed by the hedger, then authorization will be necessary. In addition, if the counterparty is one that the firm has not previously used, appropriate credit approvals will have to be obtained. The procedure for placing orders should be specified, with the hedger confirming the additional signatures necessary for entering a trade or modifying a position. In many firms order sizes above a predefined threshold will require additional signatures or authorization.

It might be prudent to *paper trade* for a period of time, that is, establish a hypothetical hedge and monitor its behavior relative to the unhedged position over a number of days or weeks, in order to ensure the effectiveness of the proposed strategy. The risk manager may ask for *back-testing* to be performed, if feasible, to demonstrate how the proposed hedge would likely have behaved in the past. In some cases, however, it may prove difficult or impossible to obtain historical prices, correlations, volatilities, and other inputs, and so they must instead be estimated, decreasing the usefulness and reliability of the back-testing exercise.

Execution

Once all internal approvals have been obtained, the next step is to establish any required accounts with the hedge provider or relevant exchanges and to execute the hedge. It may also be necessary to engage a pricing service to provide independent marks on the hedge positions as it is not desirable to be dependent on the hedge provider for price quotes on the hedging instrument. In 1994, Gibson Greeting Cards discovered to its cost the perils of depending on a hedge provider for price quotes, when it discovered that it had been misled by its counterparty, Bankers Trust, who had significantly understated the extent of Gibson's losses from derivative transactions. A comprehensive discussion of this case can be found in Chew (1996).

In many cases, risk management will require the manager to enter into a partial hedge for a period of time, to "test the waters," before deploying the entire hedge. This is particularly likely to be required if the hedge position is complex, that is, something other than a plain vanilla position, if the solution involves delta-hedging, or if historical prices required for back-testing were not available. The drawback of this approach, of course, is the consequent delay in establishing the entire hedge, during which some risk exposure will continue to be held while the testing is underway. In other situations, a

manager may deploy a hedge in anticipation of a risk exposure that has not yet arisen. Such *pre-hedging* may be appropriate when liquidity is poor and the execution of a hedge may take time, but is more common with, for example, agricultural commodities, where a farmer may choose to hedge the future sales price of a wheat harvest that has not yet been reaped. Pre-hedging exposes the hedger to performance risk, of course, with the farmer being required to deliver physical wheat (or square, i.e., close out for cash, the hedge) even if his crop is ruined.

The reader is also cautioned that, particularly with illiquid and OTC instruments, the bid-ask spread may result in a significant P/L impact when the trade is executed.

Monitoring

Once the hedge has been executed, the manager will need to scrutinize the daily P/L reports to compare the daily performance of the assets with that of the hedge in order to gauge its efficacy. Some positions may have nuances that become evident to an inexperienced hedger only with the passage of time. It might be necessary to supplement an equity option position with a volatility swap, for instance, in order to neutralize the unintended market effects of the hedge due to changes in implied volatility. In other cases, the passage of time may reveal that some assumptions in the model may have been unduly or insufficiently conservative. For example, a firm holding a portfolio of fixed-rate auto loans that it is financing with short-term, floating-rate debt may enter into an interest rate swap to hedge its interest rate risk. In this case, a swap with a notional balance that declines over time, known as a *balance guarantee swap*, might be an appropriate instrument to use to account for the fact that borrowers will repay principal over time. The declining balance schedule would be determined at deal inception by estimating the likely behavior of borrowers over time. If the observed prepayment rate is faster or slower than initially modeled, that is, if the customers chose to repay their loans earlier or later than originally anticipated, then over time the notional schedule of the swap will no longer be in harmony with the outstanding balance of the loans and so the swap will need to be rebalanced. When using nonlinear instruments, for example, or if engaged in a delta-hedging strategy, the hedge position will need to be rebalanced periodically, sometimes daily or even multiple times per day, in volatile markets, depending on risk limits.

If there is considerable basis risk between the position and the hedge, it can cause additional P/L volatility on an inter-day basis. For example, a U.S. investor hedging an investment in European mutual funds by entering into short positions in the FTSE Group (FTSE) 100 and the Deutscher Aktien

Index (DAX) 30 indices may find that the price movements of the hedge do not, on a day-to-day basis, offset the price movements of the mutual funds. Over several days, however, a movement in the hedge is likely to substantially offset a movement in the assets. If not, the hedge is not correctly specified, perhaps as the result of a change in correlation, and further analysis is required to determine if the hedge should be adjusted, or restructured.

Accounting Considerations

It is necessary when structuring a hedge and obtaining the necessary approvals for a manager to have discussions with his firm's financial reporting and accounting policy groups. The form of the hedge that is chosen may have unintended, and undesirable, implications for financial reporting, and it may be preferable to employ an alternative hedging strategy.

Financial securities may be classified in one of three reporting categories that are defined in FAS 115, "Accounting for Certain Investments in Debt and Equity Securities":

1. *Hold To Maturity* securities are debt securities that the company has both the positive intent and ability to hold to maturity. These are reported at amortized cost.
2. *Trading securities* are debt and equity securities that are held principally with the assumption of being sold in the near term for short-term profit. These are typically held by financial institutions and are generally traded actively and frequently. They are reported at fair value, with unrealized gains and losses passing through the income statement.
3. *Available For Sale* securities include all other debt and equity securities and are reported at fair value. They are generally held by nonfinancial companies, and since they are intended for longer-term returns, unrealized gains and losses are excluded from earnings and reported in a separate component of shareholders' equity.

GAAP accounting requires firms to report the market value of derivative instruments on their balance sheets. If a firm chooses to hedge securities that are classified as hold to maturity, in its effort to reduce risk, the firm may actually increase P/L volatility significantly, due to the mismatch in the accounting treatment of derivatives and the assets or liabilities that they are hedging.

Firms may choose to adopt *hedge accounting*, defined in FAS 133, "Accounting for Derivative Instruments and Hedging Activities." To qualify for hedge accounting, a hedge must, both at inception and in subsequent periods, be highly effective in offsetting changes in the fair values of, or cash

flows attributable to, the hedged asset or liability for the life of the hedge. A hedge is generally considered to be “highly effective” if the change in value of the asset or liability that it is hedging is in the range of 80 to 120% of the change in the value of the hedge.

When designing complex finance transactions and structured products that incorporate, for the purpose of hedging, *embedded derivatives*, such as interest rate caps or swaps, or equity put options, and that are intended to be accounted for using hedge accounting, the structurer must ensure that the transaction conforms to the rules regarding hedge effectiveness. If the hedge instrument is an interest-rate swap, FAS 133 provides a shortcut to ensuring hedge effectiveness by requiring that the notional amount of the swap matches the principal amount of the interest-bearing asset or liability.

The manager should always obtain professional guidance on the tax and reporting implications of a given hedge strategy.

CHAPTER 7

Constructing a Hedge

If I learned all this so slowly it was because I learned by my mistakes, and some time always elapses between making a mistake and realizing it, and more time between realizing it and exactly determining it.

Edwin Lefèvre, *Reminiscences of a Stock Operator*

7.1 AN IDEAL HEDGE

In seeking to reduce risk, one approach is to seek to minimize risk concentration through diversification, that is, by capitalizing on a lack of correlation in assets, while another is to make use of correlation, through hedging. An ideal hedge is one for which the change in value as a result of a given change in market conditions is such that it exactly offsets the change in value of the position to be hedged, that is, the changes in price of the hedge are perfectly inversely correlated with the price changes of the unhedged position. There are two steps in determining the characteristics of this ideal hedge, if it exists.

1. Identify the combination of instruments or contracts whose aggregate change in price as a result of changing market conditions is inversely correlated with that of the position. That is, determine what hedge position has the opposite risk profile to the position we wish to hedge.
2. Determine the number of those instruments or contracts required such that the change in value of the hedge exactly equals and offsets the change in value of the position we are seeking to hedge. That is, find out how big our hedge needs to be.

7.2 A SAMPLE HEDGE

Consider an oil distributor that has just entered into a contract today, which let us say is January 1, to deliver heating oil to his customers next winter at a price that is the lower of the prevailing price on December 10, or \$2.35 a gallon. Let us suppose further that the total volume of oil that the company has undertaken to deliver is 420,000 gallons, that the current spot price is \$2.20 per gallon, and that delivery is to take place on the pricing date. The distributor can purchase oil in the open market today, at \$2.20, but runs the risk that the price will drop between now and December 10th. For example, if on December 10th the spot price is \$2.00, the distributor will lose $420,000 \times (\$2.20 - \$2.00)$ or \$84,000. However, if the price rises above \$2.35 by December 10th, the distributor will gain a maximum of $420,000 \times (\$2.35 - \$2.20)$ or \$63,000. As a result, the distributor is exposed to oil price risk if he purchases oil today, being unsure whether he will ultimately gain or lose on the transaction and, in addition, would incur costs associated with storing the oil between now and December.

If the distributor chooses not to purchase oil today but to wait to purchase in the open market on December 10th (assuming, for simplicity, that the oil can be delivered to his facility instantaneously), his risk profile will have changed. If on December 10th the spot price is \$2.00, the distributor can purchase at \$2.00 but is obligated to deliver at that price, so he makes no profit. However, if the price rises above \$2.35, he is obligated to purchase at that higher price and sell to his customers at \$2.35. In fact, he is exposed to potentially unlimited losses. If the price of oil on December 10th were to be \$3.00, for example, then the distributor would lose $420,000 \times (\$3.00 - \$2.35)$ or \$273,000.

To hedge this risk, the distributor can use the energy futures market to hedge his exposure by going long the December Heating Oil Futures contract, which is traded on the Chicago Mercantile Exchange. The prices for various 2010 contracts are included in the table of futures contract prices shown in Figure 7.1. It can be seen that more distant futures have a higher price than near-term futures and are in turn higher than the current spot price (not shown in the figure); this is the normal relationship that reflects the costs of financing and oil storage, generically termed *cost of carry*. The difference between the spot price and the futures price is referred to as the *basis* and can be positive or negative, depending on the characteristics of the market for the commodity in question. The basis changes with the passage of time and converges to zero on the expiration date.

We can see that the most recent trade of the December 2010 contract was at \$2.3371. By going long the contract, the distributor can enter into a long

Heating Oil Futures

Quotes | Contract Specifications | Performance Bonds/ Margins | Product Calendar | Learn More



Quotes | Time & Sales | Volume | Settlements

Trade Date: 1/9/2010

Globex Futures | Open Outcry Futures

Market Data is delayed at least 10 minutes

View Spread Quotes

Turn Auto-refresh ☒ About this Report

Month	Charts	Last	Change	Prior Settle	Open	High	Low	Volume	Hi/Lo Limit	Updated
Feb 2010 [OPT]		2.2047	+0.0211	2.1836	2.1864	2.2081	2.1587	40,646	2.4338 1.9338	4:00:52 PM CST 1/8/2010
Mar 2010 [OPT]		2.2115	+0.0169	2.1946	2.1951	2.2145	2.1690	19,357	2.4446 1.9446	4:01:22 PM CST 1/8/2010
Apr 2010 [OPT]		2.2143	+0.0168	2.1975	2.1925	2.2150	2.1719	9,678	2.4475 1.9475	4:00:52 PM CST 1/8/2010
May 2010 [OPT]		2.2157	+0.0142	2.2015	2.2018	2.2173	2.1855	8,438	2.4515 1.9574	4:00:22 PM CST 1/8/2010
Jun 2010 [OPT]		2.2219	+0.0145	2.2074	2.2013	2.2244	2.1820	7,080	2.4574 1.9574	4:00:52 PM CST 1/8/2010
Jul 2010 [OPT]		2.2295	+0.0083	2.2212	2.2090	2.2350	2.2090	1,646	2.4712 1.9712	4:00:52 PM CST 1/8/2010
Aug 2010 [OPT]		2.2462	+0.0070	2.2392	2.2304	2.2531	2.2235	577	2.4892 1.9892	4:00:22 PM CST 1/8/2010
Sep 2010 [OPT]		2.2604	+0.0018	2.2622	2.2545	2.2620	2.2500	806	2.5122 2.0122	3:59:52 PM CST 1/8/2010
Oct 2010 [OPT]		2.2911	+0.0039	2.2872	2.2779	2.2911	2.2750	484	2.5372 2.0372	3:54:52 PM CST 1/8/2010
Nov 2010 [OPT]		2.2990	+0.0122	2.3112	2.2990	2.2990	2.2990	149	2.5612 2.0612	3:59:52 PM CST 1/8/2010
Dec 2010 [OPT]		2.3371	+0.0019	2.3352	2.3264	2.3412	2.3175 a	1,623	2.5852 2.0852	4:01:22 PM CST 1/8/2010
Jan 2011 [OPT]		2.3614	+0.0030	2.3584	2.3614	2.3614	2.3614	198	2.6084 2.1084	3:54:52 PM CST 1/8/2010

FIGURE 7.1 Prices of various Heating Oil Futures contracts.

Source: www.cmegroup.com.

position that obligates him to take delivery of heating oil in December. The futures contract is an obligation on both the purchaser and seller to transact at a guaranteed price of \$2.3371, irrespective of where the spot price is, on that date. Since the distributor has contracted to deliver to his customers at a price of \$2.35, he will be locking in a profit of $420,000 \times (\$2.35 - \$2.3371)$ or \$5,418.

Referring to our two criteria for sizing a hedge, the distributor will see upon investigation that the terms of the futures contract are for delivery on December 10, 2010, which is the same date that he has contracted to deliver to his customers, and the grade of oil is that which he has undertaken to deliver. So this contract is an ideal hedging instrument for his exposure. The second criterion is to determine the number of contracts required. The specifications of the heating oil futures contract state that one contract relates to a quantity of 42,000 gallons. Since the distributor needs to hedge 420,000 gallons, he will need to go long 10 contracts.

Imperfect Hedges

In practice, however, when hedging an exposure, the terms of a futures contract are unlikely to coincide exactly with the hedger's requirements. For example, the distributor might require delivery on December 1st, and his need might be for diesel fuel rather than heating oil. In such a case, using the December Heating Oil futures contract to hedge would expose the distributor to both timing risk—also known as calendar risk—and basis risk. When considering the basis risk introduced by the hedge, the distributor would need to decide if the risk of remaining unhedged outweighed the basis risk introduced by the hedge. He could look at the long-term correlation of heating oil and diesel oil prices and make an assumption as to whether that correlation was likely to continue, or if it was likely to change due to a systemic shift in the oil markets. It should also be noted that futures contracts are indivisible, so that if the distributor wished to hedge 475,000 gallons, for example, he would have to choose between hedging 462,000 gallons (using 11 contracts) or 504,000 gallons (12 contracts). In the former case, he would be underhedged and, in the latter, overhedged.

Since the calendar risk in our example is relatively small, the distributor may decide that it is inconsequential and, if comfortable with the terms of the proposed hedge, may enter into the contract, and rather than wait for delivery on December 10th, could close the position out on December 1st, for cash, and purchase oil in the spot market. In fact, most traders square, that is, close out, their positions before expiry. It should be noted that the profit or loss on the distributor's position will depend on the price at which he entered the futures contract. Futures prices vary over time, and, as an example, a chart of historical price movements for the December 2010 Heating Oil Futures contract can be seen in Figure 7.2.

Margin

Futures, unlike securities, are not purchased or sold. Rather, like forwards, they are contractual obligations to purchase or sell, that is, to accept or to deliver, some defined asset in the future, at a fixed price on a fixed day. The obligation exists until expiry, at which time it must be fulfilled, or until the position is squared, or closed out, by entering into an offsetting position. One important characteristic of trading in the futures market is the practice of *margining*, which is the requirement for participants in futures contracts to post collateral as a performance bond so that any failure on their part to perform on a contract would not result in a loss to the clearing house, which fulfills the role of counterparty to both the long and the short. The *initial margin* is that which must be posted to initiate a position, and it is



FIGURE 7.2 Historical prices of the December 2010 Heating Oil Futures contract. Source: www.cmegroup.com.

determined by the exchange or clearing house. Since initial margin is only a fraction of the value of the assets controlled by the futures contract, futures contracts can be considered to be highly levered. Every day, at the close of business, the clearinghouse will mark to market each open position, with the *settlement price* being the last trading price. The *maintenance margin* is a threshold level, for example, 75% of the initial margin, determined by the exchange. As a position is marked to market, gains and losses on the position are netted against the margin account. If the value of the margin account drops below the maintenance margin, then a *margin call* is made, and the account holder must post *variation margin* to replenish the margin account. Failure to meet a margin call will result in the position being unwound, or closed out, by the exchange. For example, continuing our scenario above, if the price of the futures contract were to fall to \$2.00 by expiry, for instance, and there were no margin requirements, then a failure of the distributor at expiry to purchase at \$2.3371 would result in a loss to the clearing house, who would be obligated to pay to the other party to the contract $42,000 \times 10 \times (\$2.3371 - \$2.00)$ or \$141,582. The use of margining significantly reduces the clearing house’s credit exposure.

Unlike forward contracts, which do not require margining, the use of futures contracts can at times necessitate large margin payments. Failure to anticipate and to provide for these potential payments can lead to liquidity problems. A famous case relates to the experience of Metallgesellschaft Refining and Marketing Inc. (MG), a U.S. subsidiary of Metallgesellschaft AG, a German company. In 1992, MG entered into a marketing strategy whereby

it would undertake to deliver heating oil and gasoline to its customers at a price, fixed in 1992, every month for up to 10 years. As a result, MG had short positions in long-term forward contracts going out 10 years. To hedge this exposure, MG implemented a *stack-and-roll* hedging strategy. Stack hedging refers to a technique where the nearest and, usually, most liquid contract is used to hedge all remaining future price risk under consideration. For example, on June 1, an oil producer that agreed to deliver 5,000 barrels of oil in each of the next 12 quarters would short 5×12 September contracts (each contract representing 1,000 barrels). In August, the producer could close out the September contracts and short 5×11 December contracts, thereby rolling the “stacked hedge.” An alternative approach, utilizing a *strip hedge*, is one in which the distributor hedges a quarter’s production with contracts whose maturity matches the production date. Such a hedge would require the distributor, on June 1, shorting five contracts in each of the next 12 quarters.

By implementing a stack-and-roll hedge, MG was hedging its short position in long-term forwards with a long position in short-term futures contracts that it continually rolled forward. Because long-term oil forward prices at the time were lower than short-term futures prices—a phenomenon referred to as *backwardation* (see Figure 7.3)—as time passed, MG would lose money on the ever-shortening forward positions, but with these losses




FIGURE 7.3 In backwardation, futures prices are lower than spot prices.




FIGURE 7.4 In contango, futures prices are higher than spot prices.

being offset by gains on the hedge, the long short-term futures contracts. Unfortunately for MG, however, due largely to changes in supply and demand expectations, the market switched from backwardation to *contango*, where long-term forward prices are higher than short-term forward prices, in turn higher than spot prices (see Figure 7.4). As MG rolled its short-term futures, it incurred losses. Unfortunately, under German accounting rules, the futures were marked to market on a daily basis with the losses being recognized immediately, while the gains on the short position in long-term forwards were not recognized—since forwards are not marked to market—leading MG to book massive losses.

While MG's strategy may have effectively hedged market risk, it introduced the company to significant levels of liquidity and operational risk. The scale of its positions was enormous, representing a significant percentage of total open interest on the New York Mercantile Exchange (NYMEX). MG's positions equalled approximately 160 million barrels of oil, that is, 85 days of the entire oil production of Kuwait. Unwinding the firm's positions in a timely fashion would not have been possible without significant market impact. And the scale of these positions meant that the company's cashflow needs were enormous. In fact, liquidity became such a problem that MG was forced to unwind its positions and recognize a loss of \$1.5 billion. A detailed discussion of this case can be found at www.prmia.org/pdf/Case_Studies/MG_IIT.pdf.

7.3 STATIC AND DYNAMIC HEDGING

In the oil distributor example above, the hedging strategy employed was a *static* hedge, where a hedge is sized at inception to eliminate the risk exposure, and then left unadjusted until expiration, or until the hedge is closed out. One advantage of a static hedge is that the hedge does not need to be adjusted in response to changes in market factors such as prices, interest rates, volatility, and so on. In practice, of course, the market factors will evolve and the composition of most portfolios will change over time, with securities maturing, and assets or commodities being purchased and sold, and so, as the markets and the portfolio change over time, the hedge will need to be adjusted or redesigned.

Some risk sensitivities will change as a result of changing market prices, interest rates and volatilities. For example, if we consider an equity call option, with an exercise price of \$40, and only a few days left to maturity, the set of possible prices for the option will range from practically \$0, when the stock price is significantly below \$40, that is, the option is deeply out of the money, to a value of close to \$C, when the stock price is $40 + C$ and C is significantly greater than 0, that is, when the option is deeply in the money.

In Figure 7.5 we can see that the slope of the tangent line to a call option price curve can range from zero to one. This slope represents the




FIGURE 7.5 The slope of the tangent to a call option price line changes with the price of the underlying, and ranges from zero to one.

option's price sensitivity to changes in the price of the underlying stock and is referred to as the option's *delta*. In general, delta will range from close to zero when the option is deep out-of-the-money, to one, when the option is deep in-the-money. When the stock price is equal to the exercise price of the option, the delta is approximately 0.5. The delta of an option can be calculated using a model such as the Black-Scholes model, which is described in standard textbooks on derivatives, such as Hull (2008).

If we were to write (i.e., sell) 100 call options on a stock, for example, on the fictitious ACME Corporation, we might be concerned about our exposure to rising ACME prices. If we suppose that the current price of one ACME share is \$40 and the options are at-the-money, then the delta, δ , of the call option will be approximately 0.5. That means that the change in price of the option for a small move in the stock price (such as a \$0.01 price change) will be one-half of the size of that move. Consequently, a perfect hedge for a short position in 100 call options on ACME stock will be a long position in $100 \times \delta$ shares, or 50 shares. A price increase in ACME stock of \$0.01 would mean that a long position in 50 shares would increase in value by $50 \times \$0.01$, or \$0.50, which would exactly cancel out the loss on the short position in 100 call options, which would equal $100 \times \delta \times \0.01 , or \$0.50. By establishing a portfolio of assets and options whose deltas cancel each other out, such that the delta of the overall portfolio is zero, the position is hedged against small market moves. Such a portfolio is referred to as being *delta neutral*.

In addition to hedging options with the underlying asset—which by definition has a delta of exactly one—and vice versa, we can make use of offsetting deltas to hedge one option position with another. We can also use forward and futures contracts when delta hedging and benefit from their relatively lower transaction costs. And since delta is additive, a hedge position can be constructed using combinations of options, forwards, futures, and the underlying asset.

In our example it was easy to calculate the appropriate size of such a *delta hedge* because we knew the delta. For simple call and put options, that is, *plain vanilla* options, the delta can, as we mentioned above, be determined using a model such as the Black-Scholes. For exotic options, however, calculating the delta can be quite difficult, requiring the use of a more sophisticated model, which can expose a manager to additional model risk and make delta hedging a decidedly nontrivial task. In practice, an exotic option can often be broken down into a number of constituent traded options, so that the aggregate cashflows of the constituents match those of the exotic option. Because of the *law of one price*, the market price of the exotic option should be the same as that of the replicating option position, and its delta can be determined without the use of a model.

Delta and Gamma

An even more practical obstacle to successful delta hedging that arises with any nonlinear instrument or position, including plain vanilla options, is the fact that delta is not constant. While a delta hedge will be effective for small changes in the underlying, because option prices are nonlinear functions of the underlying spot prices linear or delta hedging will fail for large price moves, in response to which the delta can move significantly. Similarly, the deltas of complex asset portfolios can change nonlinearly. As a result, as the delta changes the hedge is likely to require rebalancing, with large changes in delta warranting significant hedge rebalancing.

To delta hedge a portfolio with nonlinear exposures will require the implementation of a *dynamic* hedging strategy, one that entails continually rebalancing the hedge, often on a daily, hourly, or even more frequent, basis, with the frequency of required rebalancing being higher in volatile markets. Just as bond price responses to interest rate moves are usually calculated using both the first and second derivatives, that is, duration and convexity, similarly with nonlinear instruments, such as options, we need to make use of both delta, the first derivative, and *gamma*, the second derivative of the instrument's price with respect to the price of the underlying, to more accurately anticipate price changes. The gamma, defined and discussed in section 4.4, measures the change in delta as the price of the underlying asset changes, and should be taken into account when constructing a hedge. We can see in Figure 7.5 how the delta of a call option changes in response to changes in the price of the underlying, and it can be shown that the deltas of call and put options are most sensitive to changes in the value of the underlying when the options are near the money.

One approach to reducing the required frequency of rebalancing is to reduce the price sensitivity of delta by reducing gamma. While a dynamically hedged portfolio that uses *gamma hedging* will still require rebalancing, its degree and frequency will be less than for a purely delta-hedged portfolio. A perfectly hedged portfolio will have zero net delta and gamma, although in practice it is only possible to eliminate gamma by using an offsetting option position for every source of gamma in the portfolio. A practical implementation of a gamma-neutral strategy, therefore, sets as its goal keeping the gamma at an acceptably low and manageable level, rather than functionally at zero, particularly where a manager wishes to retain some upside potential. A gamma neutral strategy leaves the portfolio nearly delta neutral even for moderate changes in the price of the underlying.

Effective Delta Hedging

Even for a low-gamma portfolio, maintaining delta neutrality will require frequent trading of small positions and, therefore, incur high transaction

costs—a continuously rebalanced portfolio (a purely theoretically construct, of course) would pulverize itself through infinite frictional losses.

A practical hedge rebalancing process that reduces the amount of trading and does not necessarily require gamma neutrality (recall that by itself positive gamma is good) is to decide in advance at what level of price moves the manager will adjust the hedge and to size the adjustment using an effective equivalent variation investments (EVI) that would *on average* pertain until the next rebalancing. The trigger price moves can be different for up and down changes (Δp^+ and Δp^- respectively) and can be based on percentage price moves, maximum tolerable exposure (relative delta), or risk-return analysis (e.g., given a positive delta, the manager might be willing to allow the price to increase by up to 10% before rebalancing but would stanch the loss after a 2% drop). Alternatively, the manager might choose triggers Δp^+ and Δp^- such that, based on historical price volatility, a rebalancing would be required on average once a week, say, or once a month. Whatever the rationale behind determining trigger price moves, the basic idea is the same: Allow the portfolio to evolve with the markets for a while and let it build up some exposure, within acceptable limits, before trimming the delta.

The *effective delta*, determining the effective EVI via equation (4.13), should be calculated with the help of *finite* value changes ΔV^+ and ΔV^- for the up-and-down price moves respectively (i.e., one should use *secant derivatives*). This will ensure that the effective delta reflects the exact model valuation, that is, incorporates the contributions of gamma, speed, and all the higher price sensitivities. Moreover, since large price moves are inevitably accompanied by significant changes in other factors, most notably price volatility, thus ΔV^+ and ΔV^- should be calculated by allowing for properly adjusted volatility co-movements. (Co-movements of other important factors, especially interest rates, as well as the effect of time decay, ideally should also be included in the analysis.) The appropriate co-movement size can be determined from historical observations, as discussed in section 4.3, or, more commonly and usually at least as adequately, from traders' and managers' recorded, perceived, or intuited *volatility map* (Taleb, 1997), that is, the volatility level appropriate, in the current environment, for each price level. Since the up-and-down trigger price moves can be different, as well as the corresponding value changes, the effective delta used to size the hedge is calculated as the average of up and down effective deltas:

$$\delta_{\text{eff}} = \frac{1}{2} \left[\frac{\Delta V^-}{\Delta p^-} + \frac{\Delta V^+}{\Delta p^+} \right]. \quad (7.1)$$

Effective delta (sometimes also called *modified* or *shadow* delta) is conceptually similar to the scenario grid, introduced in section 4.2, and to effective duration (equations (4.52) and (4.53)) and other effective interest

rate sensitivities, described in the same section, in that all these measures strive to reflect the full complexity of the market system and recognize the interplay of various market factors and the analyzed asset. The effective delta technique is known on trading desks as *shadow gamma* (Taleb, 1997) and is also discussed by Calamos (2003), who uses the term “modified delta.”

Delta Riding

An additional potential benefit of effective delta hedging of a low-delta portfolio is *delta riding*, or *gamma climbing*. If one starts with an exactly delta-neutral portfolio that has a positive gamma, then, ignoring time decay (theta), a price move in either direction will generate a gain. Waiting before rebalancing long enough for a meaningful gain to accumulate, specifically, big enough to overcome trading costs and theta, the manager can lock in the profit by rebalancing the hedge, that is, making the portfolio’s delta zero again. The process is illustrated in Figure 7.6.

Some readers may have noticed the resemblance of delta riding to gamma trading of options and other derivative instruments. Indeed, the underlying principles are similar. In the case of effective delta hedging, the




FIGURE 7.6 Delta riding in effective delta-hedging process. Effective value change incorporates the impact of nonlinearities (higher-order sensitivities) and co-moves of volatility and other relevant market factors. The portfolio is delta-neutral when the underlying price is p_0 and is rebalanced when the price subsequently attains levels p_1 and p_2 , locking in gains ΔV_1 and ΔV_2 , respectively. The manager should superpose the effects of theta and transaction costs.

potential for gain is not a goal in itself but arises as a welcome by-product of a churn-reducing hedging regime. As is the case with gamma trading, the potential for profitable delta riding is highest in times of high market volatility.

It should be clear from the preceding discussion that dynamic hedging strategies can have significantly higher transaction costs than statically hedged portfolios due to the need for continual rebalancing of the hedge position. A compensating attraction of delta hedging, compared with a static *linear* hedge, is the potential for capturing some market upside via delta riding. The corresponding risk, of course, is that with large market moves the hedge may prove inadequate, although the use of gamma hedging can mitigate this danger. Effective delta hedging can both reduce frictional losses and increase the upside capturing potential.

In practice, index futures contracts are usually used when delta hedging in preference to stock or option positions, because of their lower relative transaction costs, although if the position to be hedged consists of a small number of equity positions, then index futures may not be appropriate. We discuss the tradeoff between costs and upside in section 7.4.

When delta hedging or gamma hedging, the hedger should remember that the price of hedging instruments, especially nonlinear ones, can be sensitive to factors other than the price of the underlying asset: the prevailing level of interest rates, the sensitivity to which is measured by the *rho*, the *duration*, and other interest rate measures, discussed in sections 4.4 and 4.5; the passage of time, quantified by the *theta*; and most crucially changes in volatility, the price response to which is the *vega*. A delta naïve delta-hedging strategy that ignored these sensitivities could expose a manager to significant basis risk.

The *hedge ratio* is the ratio of the change in price of the hedging instrument to the change in price of the asset (or liability), and so is used to determine the appropriate size of a delta hedge. For a call option, therefore, the hedge ratio is equal to the option delta. In our discussion of ACME stock above, we saw that to delta hedge a short position in 100 call options we would need to go long 50 shares of stock; equally, to delta hedge a long position of 50 shares of stock we could enter into a hedge by shorting 100 call options, or going long 100 puts. In this case we are hedging by making use of the price sensitivity of the options, not by any expectation of exercising the puts.

7.4 PROXY HEDGING

In the hedging examples we have considered thus far, the exposure to be hedged and the hedging instruments have been readily identifiable. In the

real world, of course, there can be many situations where the actual risk exposures cannot be directly observed or measured, or where the perfect hedging instruments may not be available. In these cases, the use of *proxy hedging* is warranted. Consider, for example, a manager who has invested in a portfolio consisting of positions in a variety of mutual funds, the values of which depend on the prices of a large number of domestic and international bonds and equities. Rather than attempting to understand the risk of the portfolio in terms of the market behavior of the holdings of the constituent mutual funds, which becomes a nontrivial task quickly as the number of funds increases, the manager may instead choose to create a hypothetical, *replicated* portfolio of a small number of market indices, constructed such that this hypothetical portfolio exhibits the same broad risk sensitivities as the original portfolio. The process of identifying the market indices used to model the portfolio, and their respective weights, can be effected through the use of a multiple regression technique such as a *returns-based style analysis* (Sharpe, 1992). By regressing the historical returns of the portfolio to the set of market indices, a set of index weights can be determined such that a hypothetical investment in the indices in proportions equal to the regression betas would have produced returns for this proxy portfolio similar to that of the portfolio under consideration.

The sensitivity of the proxy portfolio to changes in the levels of each of the indices, that is, the deltas and gammas of the portfolio, could then be calculated using the valuation model, as described in Chapter 4. Table 7.1 shows the sensitivity of such a hypothetical portfolio to a 1% change in various market indices. For example, if the S&P 500 were to decrease by 1%, the portfolio would decrease in value by approximately \$1.28 million. This exposure could be hedged in a number of ways, such as shorting \$128 million worth of S&P 500 index futures by entering the short side of a futures contract. A decline in the S&P 500 of 1% would increase the value of the short position by \$1.28 million, nearly exactly offsetting the decrease in the value of the portfolio brought about by the decline in the index. The exposures to the other market indices could be hedged in a similar manner. These exposures to market indices represent a set of equivalent variation investments (EVI).

It is likely, of course, that in practice such a portfolio would exhibit some gamma, so that the prices of a futures hedge and the portfolio would diverge for all but small movements in an index. Thus, in practice the hedge position would ideally be adjusted to make the position gamma neutral.

In Section 7.1 we described the two key steps in structuring a hedge, the first being the identification of the correct instruments to use in a hedging strategy, and the second being the determination of how the hedge should be sized. While the exposure hedge ratio indicates how much of the total

TABLE 7.1 Sensitivity of the valuation of a hypothetical portfolio to a 1% change in various market indices. To effectively hedge such a portfolio, a combination of positions with equal but opposite sensitivities would be required. Note that for larger market moves, gamma must be taken into account.

Index	EVI per 1%
S&P 500	1,281,812
Nasdaq 100	1,206,776
Russell 2000	1,064,797
MSCI Canada	215,495
MSCI Europe	787,465
MSCI Far East	-312,714
LB Government/Credit	162,508
ML High Yield	512,652
3 Month Treasury Bill	217,943
<i>Total</i>	<i>5,136,736</i>

exposure at risk that we choose to hedge, the actual hedge ratio may not be immediately obvious. This is particularly true when there is basis risk, that is, when the characteristics of the portfolio and the hedging instrument are different. Basis risk is unavoidable when proxy hedging and can be introduced both by portfolio replication and by the choice of hedging instruments. We have seen that when hedging a large portfolio of equities, or mutual funds, for example, hypothetical proxy portfolios are often constructed, based on market risk exposures, or EVIs, identified through the use of style analysis. Very often the risk matching will not be perfect, and the coefficient of determination, R^2 , will be less than 100% (see Appendix B for a discussion of statistical topics), reminding us of what we may refer to as the third rule of hedging, that is, that the only perfect hedge is in a Japanese garden. When sizing an effective hedge, the objective should be to ensure that the risk sensitivity of the proposed hedge matches as closely as possible that of the position to be hedged and should be adjusted to reflect any imperfect correlation between the position and the hedge instrument.

For example, consider an investor who wishes to hedge the market risk exposure of an equity portfolio consisting of \$100 million split equally between two U.S. mutual funds. If the investment objectives of the two mutual funds suggest that one is likely to invest solely in large capitalization U.S. equities, while the other is likely to invest solely in U.S. technology companies, then performing a style analysis on the portfolio by regressing its historical returns against those of a set of U.S. broad market indices would very likely

result in a proxy portfolio being identified that consisted of a 50% exposure to the and a 50% exposure to the NASDAQ 100. The R^2 of the regression, however, will almost certainly be less than 100%, since the funds will no doubt hold some of their assets in cash, they might make use of leverage, and they might not invest in the same equities, in the same weights, as the two indices. The behavior of the actual portfolio, therefore, will differ from that of the proxy, and a hedge designed to exactly offset the exposure of the proxy will fail to perfectly hedge the actual portfolio—there will be basis risk. As a result, hedging the proxy portfolio, by shorting \$50 million each of S&P 500 and NASDAQ 100 futures, for example, would result in an imperfect hedge. In this example, the hedged position would be overhedged, so that the hedged portfolio would tend to lose money if markets appreciated, and would tend to gain if markets depreciated. Consequently, the hedge must be reduced, that is, adjusted to account for the degree of correlation between the portfolio and the hedge, and the respective volatilities of both the hedge instrument and the underlying asset, with the objective being to minimize the variation in the return of the hedged portfolio.

7.5 PROTECTION VERSUS UPSIDE

The discussion of the hedging process above suggests that the successful implementation of an effective hedging strategy can involve a great deal of time and effort. Even a relatively unsophisticated static hedge can require a significant investment of resources in researching and identifying one or more appropriate hedge instruments, creating the required accounts, setting up pricing feeds, and so on. In addition to these operational costs, insurance comes at a cost, and so there can also be significant opportunity costs that result from foregoing potential profits. For example, in the case of a farmer who chooses to hedge his corn crop for delivery next summer by locking in his sale price he is ruling out the possibility of profiting from an increase in the price of corn. While this may very well be the appropriate strategy for a relatively unsophisticated enterprise, a large, well-capitalized national food company, however, might form the considered opinion on the basis of significant research and macroeconomic analysis that the likelihood of an increase in corn prices significantly outweighs the probability of a decline in prices, and that on a risk-adjusted basis remaining unhedged is a reasonable strategy. In this case, the company is prepared to retain the downside risk in order to have the chance of benefiting from a price increase.

In practice, the relevant question to be answered is not whether to hedge, but rather, in the context of the enterprise's willingness and ability to take risk, what is the appropriate degree of exposure to retain? The answer to

this question depends on the risk appetite of the enterprise, and the context. For example, a manager entrusted with the task of managing money for an elderly retiree is likely to adopt a more conservative hedging strategy than when managing assets for a 35-year-old who is still working, while a self-made entrepreneur who has a proclivity for risk taking might be more aggressive than someone who has inherited wealth. The ability to take risk will differ among individuals according to their net worth, earning potential, and age, while their willingness to take risk will also likely vary with age. For an enterprise, if downside market volatility exposes the firm to the risk of bankruptcy, then the decision to fully hedge and forego potential profits seems a rational decision. For a fully funded pension plan, however, hedging a position in the equity markets using a dynamic hedging strategy that retains the potential for some upside participation seems perfectly justified, as would an alternative strategy of funding disaster insurance, by purchasing out of the money puts for example.

This leads to the key questions of whether the risk exposure should be hedged statically or dynamically, and for how long. As we have seen, a linear static hedge is relatively straightforward to size and execute, with the tenor of the hedge matching the tenor of the exposure and minimal monitoring required, as long as the portfolio being hedged remains unchanged. The potential for upside participation, however, is minimal if the exposure hedge ratio is 100%, that is, if the entire exposure is hedged, and if an effective hedge is implemented using linear instruments such as forwards or futures. This type of hedge is usually *unfunded* in the sense that no capital is necessary to initiate the hedge (other than the deposit of the initial margin required when using futures contracts). Alternatively, the hedger can choose to make use of a *funded* hedge, entering into an insurance-like contract by purchasing nonlinear instruments such as, for example, put options. The use of put options allows a hedger to protect against adverse market movements while retaining the potential to profit from favorable market moves. It should be noted that with plain vanilla put options, gains from market appreciation are not protected and will be lost if the appreciation proves temporary, reducing the attractiveness of the strategy. With funded strategies, the protection against market risk and any appreciation are at the expense of the hedge premia that have been paid.

Portfolios with nonlinear exposures, such as option positions and levered positions in equity, for instance, provide the potential for significant upside appreciation, particularly if they are dynamically hedged. Consider, for example, a mutual fund manager that wishes to hedge the exposure of its management fees to market risk. Since management fees are a function of the net asset value (NAV) of the fund, market price volatility will have a significant impact on the present value of the fee stream, which will have

positive gamma exposure due to the compounding effects of price growth. This delta and gamma exposure is best hedged by a dynamic hedging technique. Using a valuation model, a set of equivalent variation investments (EVIs) in market indices can be determined and can be used to size an appropriate short futures hedge. If the NAV of the fund increases, the short position will expose the manager to an increased liability as the underlying indices increase. If the NAV of the fund decreases, the short position will increase in value. In both scenarios, the change in value of the short futures position will offset the change in value of the fee stream, assuming the hedge is appropriately sized. A change in NAV value will also lead to a change in the fund's EVIs, since the delta and gamma will change, thereby requiring a rebalancing of the dynamic hedge.

One significant disadvantage of dynamic hedging, particularly in volatile markets, is that the constant rebalancing of the hedge leads to transaction costs that can be considerable. Indeed, the process of dynamic hedging increases volatility, tending to increase the cost of the strategy itself in a positive feedback loop. To mitigate transaction costs a strategy that employs a combination of static and dynamic hedging can prove effective in balancing the potential for upside while containing transaction costs. A static hedge can be used to provide catastrophe insurance, for example, limiting losses to no more than 20% of the portfolio value (plus option premiums) by purchasing puts with an exercise price of 80% of initial market levels, while dynamic hedging can be used to hedge 20% of the portfolio. If the portfolio were to experience a sustained increase in value, then the catastrophe insurance could be ratcheted up, locking in gains over time. The decision to choose static or dynamic hedging will be influenced by the level of market volatility. Periods of high volatility tend to favor static hedging, although, consistent with the supply-demand argument, in such times nonlinear instruments, such as options, will tend to be more expensive compared to periods of lower volatility. Linear instruments allow good delta-hedging and are, as a rule, relatively inexpensive. Nonlinear ones, however, often have positive gamma, allowing one to keep some upside but typically have greater delta mismatch and thus lead to higher P/L volatility.

A second, potentially more serious, disadvantage of certain kinds of dynamic hedging relates to the need to rebalance the portfolio in volatile markets. Just as stop loss limits can trigger the execution of an order, but do not guarantee the ultimate level at which the trade will take place, similarly dynamic hedging can expose a manager to the same uncertainty. When hedging a position with plain vanilla put options, for example, there is a guaranteed floor for the position, which is the exercise price of the option. With dynamic hedging strategies such as option replication and portfolio insurance, there is a need to sell in falling markets. Practitioners of portfolio

insurance learned to their cost during the 1987 stock market crash that the desire to sell and the ability to sell don't always coincide. For a thorough examination of this topic see Jacobs (1999).

7.6 BASIS RISK

In the preceding sections, we discussed how basis risk can occur when hedging due to a difference in characteristics between the position to be hedged and the instruments used to hedge. Those differences can be the result of an imperfect correlation between the position and the hedge, a difference in timing, or the use of proxy hedging. In Chapter 2, we suggested that risk is often measured in relative terms, with the performance of an investment portfolio being usually measured against a reference index or benchmark, usually a standard market index such as the S&P 500 or the Barclays Capital U.S. Government/Credit Index, although occasionally the benchmark can be a custom-tailored index such as a mixture of standard indices (e.g., 60% S&P 500 and 40% Barclays Capital U.S. Government/Credit Index for a balanced fund) or a custom liability index. The latter can be appropriate for a pension fund that has a series of obligations in the future. The Financial Accounting Standards Board (FASB) in 1985 issued Statement of Financial Accounting Standards No. 87, "Employers' Accounting for Pensions" (FAS 87), which recommended that, in order to value its pension obligations, a firm construct a portfolio of high-quality zero-coupon bonds whose maturity dates and amounts match the timing and size of its future pension benefit payments. Managing its pension portfolio against any other benchmark could result in the eventual underfunding of the firm's pension plan, even if it were fully funded originally and the plan assets outperformed the benchmark, due to the use of an inappropriate bogey. Deviations in performance from the benchmark index, which are due in large part to basis risk, can be quantified by a measure known as *tracking error*, the standard deviation of the difference between the returns of the portfolio and of the benchmark:

$$\text{Tracking error} = \sqrt{\frac{\sum_{t=1}^T (\alpha_t - \bar{\alpha})^2}{T - 1}} \quad (7.2)$$

where α_t is the excess return of the portfolio over the benchmark (also referred to as alpha) in period t and $\bar{\alpha}$ is the portfolio average excess return.

The effectiveness of a hedge can be quantified in the same way, with a larger tracking error indicating a less effective hedge. For example, when hedging a basket of stocks using S&P 500 futures, the tracking error is equal to the standard deviation of the stock basket's excess return over the index.

If the tracking error is small compared to the volatilities of the asset and the S&P 500, then a hedge constructed using index futures should prove effective. In such a case, the minimum variance hedge ratio would be close to one, since the coefficient of correlation ρ would also be close to one, signifying that the two return series move in lockstep, and the volatilities of the asset and the hedge should be similar. When comparing return series, the choice of sample history period and duration is important. If correlations have changed over time, it is important to understand why, in order to help assess the likelihood of correlations changing in the future. With a mutual fund, for example, the fund's investment objectives, strategy, or benchmark might have changed, leading to a change in investment style. In addition, the choice of return period (hours, days, weeks, months) can influence the measured degree of correlation between an asset and an index, particularly if there is relatively low liquidity for an asset. When hedging a position that is to be marked to market on a daily basis, for example, daily returns are more likely to be appropriate for use in sizing the hedge than hourly returns.

In section 7.4 we mentioned the use of returns-based-style analysis to identify a suitable proxy for hedging. The same methodology can be used to identify changes over time in the composition of a portfolio. For example, an investor who is seeking to hedge the market risk of a mutual fund might select subsets of the return history of the portfolio, such as the returns of three 36-month periods, ending one month ago, 25 months ago, and 49 months ago, respectively. Using these time slices, he could perform a style analysis on each of the three periods of history, the result being three sets of exposure factors, or EVIs. If the risk exposures of the fund are determined to have changed significantly over time, it is likely to be due to the fund manager having altered the fund's investment strategy. The investor would need to be satisfied that no more changes in strategy are likely to occur during his future time horizon if he is to be confident that a hedge based on the current EVIs will prove effective. An alternative approach to determining the style of a fund is to analyze the list of the fund's reported holdings, and their respective weights, categorizing the holdings in their appropriate market sectors. This approach is particularly useful for those funds or pools for which the use of style analysis can prove problematic, such as hedge funds that use long/short strategies, for example.

When hedging the market risk of funds or pools whose composition is likely to change over time, a dynamic hedging strategy is appropriate. In this case, as the investor periodically calculates updated exposure weights, or EVIs, he can alter the composition of the hedge in response to changes that have occurred in the weights. The risk, of course, is that of adverse market moves in the time period between the manager's change in investment style and fund composition, and the identification of that change by the investor

and his subsequent *beta adjustment*. Any delay between the two events will result in basis risk while the hedge is less than optimally structured. The beta adjustment process makes use of a layered hedge approach, with the adjustment effected through the addition or removal of incremental hedges, rather than the replacement of the entire hedge.

7.7 UNINTENDED CONSEQUENCES

We have previously pointed out that basis risk can be introduced as a result of an imperfect hedge, using the example of an oil producer who wished to hedge the production of oil in Dubai using a hedge based on a different grade of oil. Because it was not feasible to hedge the exposure directly, basis risk was introduced. While basis risk in this example arose from an inability to enter into a perfect hedge, in other cases a seemingly perfect hedge can have unintended consequences. For example, the manager of a portfolio that includes nonlinear instruments such as options, or instruments with embedded optionality such as convertible bonds, may be exposed to volatility risk. For such a portfolio, a reduction in equity market volatility would reduce the value of equity options. To hedge such vega exposure, the manager could enter into a volatility swap, going short vega. If the manager was not familiar with such a hedge, he might be taken aback to discover unexplained P/L volatility in the daily mark to market reports on his hedged position. Further investigation would likely reveal to the manager that the volatility swap, while successfully reducing the vega exposure of his portfolio, had, unbeknownst to him, increased the delta of his entire position. He would, no doubt, also be surprised to learn of the negative gamma introduced by going short the swap.

It can be useful to think of gamma as adding or subtracting additional exposure to an asset. This becomes particularly clear if one expresses gamma as a dollar impact in response to a unit change in the price of the underlying. For example, if an equity position had an EVI of \$10 million and a \$1 price change gamma impact of \$1 million, then if prices were to increase \$1 the resulting exposure would be \$11 million (assuming the position was delta-neutral). Similarly, consider an oil trader with a notional position of 50 million barrels. If the gamma of his position was \$5 million barrels per \$1 price move per barrel (bbl), then a \$1 per bbl increase in price would cause the trader to effectively be long 55 million barrels, with a \$1 per barrel price decrease would reduce his position to being long 45 million barrels. Gamma, therefore, can change the profile of a position, even a hedged one, very quickly.

Basis risk can also arise from the fact that volatility (vol) is a complex characteristic, with a *volatility surface* reflecting different levels at different maturities and price levels. The shape of the vol surface will vary for different asset classes and instruments, with deep out-of-the-money equity options, for example, being priced to reflect levels of volatility usually, although not always, higher than for at-the-money options. In the crude oil markets, supply and demand considerations can lead to at-the-money options having higher volatility than deep in-the-money options. Because of the shape of the vol surface, with volatility depending on maturity and strike, incorrectly specifying the terms of a volatility swap—such as using one-year vol to hedge five-year vol—will introduce basis risk. This is a particular concern for exotic options such as barrier options, which are priced using several different vol levels.

Other unanticipated dangers can arise from the fact that portfolios made delta-neutral by the purchase of options will have a negative theta, and so will suffer from *time decay*, with the option value eroding over time, reducing the mark to market value of a hedge. In general, shorter-dated options decline faster in value than those with longer times to expiration. As a consequence, a manager intending to hedge short-term exposures may find it desirable to purchase longer-term options, and sell them when the position no longer needs to be hedged.

Because of unintended consequences such as time decay, it may be prudent, as we suggested previously, for managers proposing to enter into hedges constructed using instruments with which they have no practical experience to hedge only a portion of their total exposure while they gain confidence and experience with the new instruments. The danger, however, once again, of such caution is the risk of being less than wholly hedged during the trial period. Finally, the manager must decide if the unintended consequences of a proposed hedge outweigh the benefit provided by the hedge, and if some exposures might be best left unhedged or only partially hedged.

7.8 HEDGING CREDIT RISK

Credit risk associated with financial instruments is the risk of loss due to the default of or nonperformance by an issuer, loss due to a change in credit spreads, and loss resulting from a credit downgrade. Credit risk will also exist for banks and other institutions that have extended loans to customers, while purchasers of certain types of asset-backed securities, for example, those backed by collateral such as credit card and auto loan receivables, are also exposed to consumer credit risk. While spread risk can arise due to a

crisis of confidence in the marketplace, or a reduction in liquidity caused by, for example, a “flight to quality,” rather than a change in the likelihood of an obligor fulfilling his obligations to repay, in general credit risk arises from the fact that the interests of creditor and obligor are divergent, with an information asymmetry existing between the equity holder and the creditor. Because of limited liability, the equity holder, who in effect has a call option on any enterprise upside, often has an incentive to engage in speculative ventures that, if they fail, will adversely impact the interests of the lender, whose upside potential is capped. Recent notable examples of this phenomenon include Enron, a company that changed the focus of its business from owning gas pipelines to energy trading, and the Seagram Company, which transformed itself from a liquor distiller to an entertainment group. Many consumers, too, were less than forthcoming in the runup to the recent credit crisis when applying for credit facilities such as mortgages, with “liar loans” —those issued without the furnishing by the prospective borrower of traditional documentation verifying income and assets—being quite common.

Exposure to individual consumer credit risk is very difficult to hedge. Financial institutions such as banks that extend mortgages, credit card facilities, or auto loans to consumers may attempt to diversify their exposure by seeking to ensure that the pool of borrowers is very large and is not confined to one geographic region. Financial institutions owning receivables such as auto loan and credit card pools can make use of techniques such as *static pool analysis* to analyze the behavior of borrowers in prior pools and attempt to extrapolate how current pools are likely to perform, taking into consideration any changes in underwriting standards and macroeconomic projections. Such firms may identify a relationship between borrower defaults and interest rates, for example, and seek to hedge using interest rate derivatives. In practice, though, hedging loan portfolios is difficult. Instead, the focus for financial institutions is generally on *loss mitigation*, first by setting an interest rate on a loan that is commensurate with the estimated risk of the transaction and second, for collateralized loans such as auto or mortgage loans, seeking to ensure that the loan amount is less than the value of the collateral so that, in the event of default, the collateral may be seized and sold, in the hope that the proceeds of the sale will exceed the loan balance and the costs of asset disposition. In addition, in many cases the lender will seek additional credit support by requiring the borrower to provide a personal guarantee or a letter of credit, or, in the case of mortgages, private mortgage insurance (PMI).

Since the 1980s, lenders have increasingly turned to the sale of these consumer receivables through securitization. By selling the assets, often while recognizing a *gain on sale*, the transferor could divest itself of their credit

risk, while simultaneously freeing up its balance sheet, allowing it to focus on originating new business. This *originate to distribute* model was the undoing of a number of financial institutions engaged in extending mortgages to sub-prime borrowers—those with less than excellent credit—in the late 2000s, when the market for these pools dried up as Wall Street banks faced liquidity issues, with the originators being left holding pools of substandard loans whose quality, in many cases, had not been examined too carefully in the expectation of a quick securitization.

For most managers, credit risk exposure is not to individual consumers and small enterprises, but rather to corporate firms, in the form of loans or bonds that, in the interest of brevity, we shall refer to collectively as bonds. Managers may also be exposed to aggregated consumer credit exposure through the purchase of asset-backed securities (ABS) and, in particular, mortgage-backed securities (MBS). It should be noted that the credit risk of both of these types of security is not to the issuer, but rather to the underlying collateral supporting the transaction, upon which the bond cashflows are dependent. For example, the interest and principal due to the bondholders of a bond issued by an entity such as “CHASE CREDIT CARD MASTER TRUST” are obligations not of Chase Bank, USA, National Association, or its affiliates, but rather of the bankruptcy-remote entity to which the credit card receivables have been transferred in a *true sale*. Each of the receivables is an obligation of an individual credit card holder, who may or may not be employed, own his own home, and who may live in any of a number of states. Demographic information describing the makeup of the pool of obligors is generally provided in the bond’s *offering memorandum* and can be used by investors to determine if the pool is sufficiently diversified for their needs, and if it matches their investment criteria. To hedge the credit risk of such an issue directly is practically impossible. Instead, a bond holder could seek to reduce his risk exposure by holding only a small percentage of his total portfolio in any given issue, and if concerned about the impact of potential future interest rate increases on the ability of bond obligors to repay their credit card balances might choose to hedge that hypothesized risk using, for example, an interest rate swap.

Credit Ratings

The starting point for analysis of the credit risk is a bond’s *credit rating*, which is a measure provided by a rating agency such as Standard & Poor’s, Moody’s Investor Service, or Fitch Ratings, reflecting their opinion as to the creditworthiness of the obligation. A credit rating can relate to the issuer—an *issuer credit rating*—or it can relate to a particular bond—an *issue-specific credit rating*. Bond ratings are provided for both short-term and long-term




FIGURE 7.7 Rating agencies use a scale to denote the relative creditworthiness of a company or issue. The suffixes “+” or “-” (S&P, Fitch), or “1,” “2,” or “3” (Moody’s) may be appended to a rating to denote relative status within rating categories AA/Aa to CCC/Caa.

obligations; we shall focus on the latter. A bond’s rating will fall in a category ranging from “AAA,” for bonds of the highest credit quality, to “D,” for bonds currently in default, with a default defined as a failure to pay interest on a timely basis, and principal as it falls due, which for most sovereign and corporate bonds is at maturity.

Figure 7.7 shows the rating scales for Standard & Poor’s, Moody’s, and Fitch. While the rating methodologies of the three agencies differ, with Moody’s ratings reflecting the agency’s estimate of the bond’s expected loss, and the Standard & Poor’s and Fitch ratings reflecting the respective agency’s estimate of the probability of the bond’s default, in practice the ratings are considered comparable in the marketplace with divergent, or “split,” ratings being rare. While it may come as a surprise to many investors, it should be noted that even “AAA/Aaa”-rated bonds are considered to have a non-zero probability of default, so that the expected loss on the bond is greater than zero, meaning that the expected return of a par bond will always be less than its stated coupon.

Many financial institutions will assign their own internal ratings to borrowers and issuers. Their analysis will make use of the creditor’s annual financial statements and often will utilize a multifactor regression relationship, similar to *Altman’s Z-score*, or models such as Merton’s contingent claims model or the KMV Creditor Monitor to estimate the probability of default (PD) of a firm. Historical data on defaults can then be used to estimate the severity of loss (SOL), also known as loss given default (LGD), with the expected loss for an issuer being the product $PD \times LGD$. A *transition matrix* can also be produced, showing the probability of migration from one

rating category to another over time. For a discussion of default models and the rating process, the reader is referred to de Servigny and Renault (2004).

When evaluating the risk of a particular bond, investors should understand that rating agencies typically rate “through the cycle,” rather than adjusting their ratings to reflect business upturns or downturns. As a result, for a given rating, a bond will generally prove riskier at the end of an economic expansion than at the end of a recession. Furthermore, agencies can sometimes be slow to react to events that reduce a bond’s creditworthiness. So while the default probability of a bond can be estimated by referring to the rating agencies’ historical default tables, market participants often use changes in current market spreads for a bond for an indication of the market’s perception of the creditworthiness of an issuer relative to its peers, with wider spreads suggesting an increased likelihood of default for the bond. Spread risk and downgrade risk are influenced by supply and demand, since many institutional investors are prohibiting from investing in sub-investment grade bonds (those with a rating lower than “BBB-/Baa3”) and with some investors being restricted to even higher credit categories. As a result, as bonds are downgraded, their spreads can widen out even more than the increase in their expected default probability or expected loss would merit, as lower quality bonds typically trade at spreads that incorporate a significant liquidity premium.

Credit Protection

Investors seeking to hedge an individual corporate obligation can make use of a credit derivative such as a *credit-default swap (CDS)*, which provides insurance¹ against the default of an issuer. In return for the payment of a fee, usually expressed as a percentage of the notional value of the swap, the swap provider undertakes to make the purchaser whole in the event of a default on the reference instrument. The payment to the purchaser in the event of a loss can be either the *loss given default*, that is, the loss after any recoveries—cash settlement—or the swap provider can purchase the defaulted bond for par from the insurance purchaser—physical settlement—depending on the terms of the swap. Most CDSs are standardized and, as a result, liquid, their prices varying in the market over time as the market’s perception of the risk of the issuer changes. Investors can use changes in CDS spreads just as they use changes in bond spreads to infer an issue’s credit rating. Their liquidity make CDSs attractive, especially to firms holding illiquid bonds or loans.

¹Unlike, say, home or car insurance, the purchaser of a credit derivative need not own the reference asset.

Credit exposure to a single “name” can also be hedged using a *total return swap*, where an investor can pay the total return of a bond (including any principal loss) in return for receiving a benchmark return plus a spread. Alternatively, credit spread risk can be hedged using a *credit spread option*, where the option payout is determined by the bond’s spread relative to a benchmark, usually a U.S. Treasury bond.

The danger of instruments that provide credit protection, however, is that they do not eliminate the exposure to the credit risk of the issuer, but instead backstop it with the credit risk of the CDS writer. The American Insurance Group (AIG) crisis in 2008 demonstrated very clearly what can happen when a single CDS provider who has written protection against the default of a large number of companies sees a sharp rise in claims.

For portfolios consisting of a number of bonds or loans, it can be impractical to hedge each obligation individually using CDSs. As mentioned above, such portfolios of fixed income instruments are typically managed against a benchmark, often a broad market index such as the Barclays Capital U.S. Government/Credit Index. If the average credit rating on the portfolio securities is relatively high, such as “AA,” for example, the manager may not be unduly concerned about hedging credit risk. For portfolios with lower weighted average ratings, however, such as a portfolio of high yield bonds, managed against, for example, the BofA Merrill Lynch U.S. High Yield Constrained Index, the manager may deem it prudent to employ a hedge to manage the portfolio’s credit risk. An ideal hedge in this case might be a total return swap on the index, assuming that the portfolio’s tracking error was small. In practice, however, it may be difficult or impossible to find a counterparty to such a hedge, even at a high cost. Since swaps are a *zero-sum game*, a manager seeking to offset credit risk via such a swap would require a counterparty willing to assume the credit risk. Traders and other market makers may be constrained by their risk managers from assuming large positions. In the real world, fear and market sentiment may lead to an absence of appetite for the specific risks that a manager would wish to lay off precisely when market conditions prompt the manager to seek such a hedge.

A manager who is concerned about losses on a fixed-income portfolio and who finds it undesirable or impractical to employ a total return swap to hedge his exposure can instead choose to decompose the risk of the portfolio into pure interest rate risk and credit risk constituents and to hedge these components separately, using more liquid instruments. He can choose not to hedge the interest rate risk or, alternatively, he can hedge it effectively using, for example, an interest rate swap or a futures position. The credit risk, that is, the spread and default risks, can then be hedged in isolation. If the portfolio is small, the credit risk can be hedged using a series of credit default

swaps on each “name,” or a customized basket swap can be purchased. In the latter case, the pricing will depend on the payout formula, with the swap paying out on the default of a single “name,” for example, or, with an n th to default swap, paying out only after n companies were to default.

The manager could hedge the credit risk of a large portfolio using a liquid CDS index such as CDX or iTraxx or, for a portfolio of ABS instruments, the ABX index. The disadvantage of hedging using such an index is, of course, the basis risk it introduces. Unless there is a high degree of correlation between the issues in the portfolio and those referenced in the index, there is the potential for significant tracking error. As discussed in Chapter 2, there are significant idiosyncratic differences between companies, even within the same sector, that can lead to different relative performance. While some companies within an industry may be forced to file for bankruptcy, others may prosper.

Another potential source of basis risk arises from changing consumer behavior. For example, an investor who is exposed to consumer credit risk, by owning bonds backed by credit card receivables, for instance, may use the ABX index in an attempt to hedge his credit risk. The protection afforded by such a hedge will actually be against the risk of default on subprime mortgages, which is influenced largely (although not exclusively) by interest rates and housing prices, and so may not be an ideal hedge against the risk of the defaults by credit card holders, who may choose to make their home loan payments at the expense of paying their minimum credit card payments. Similarly, holders of mortgage debt who choose not to hedge their credit exposure, believing that homeowners will always seek to make their house payment ahead of any other obligations, may discover that consumers are opting to default on their mortgages, preferring to direct their resources to ensuring that their credit card lines are protected so that they can continue to pay for groceries, and to rent an alternative dwelling. The lesson for investors exposed to consumer credit risk is that it is difficult to anticipate how consumers will behave in the future.

It is particularly important when entering into credit risk hedges to ensure that the documentation, which is necessarily complex, specifies the conditions under which payments will be made. For example, corporate debt and bond issues can include restrictive covenants and cross-default and other provisions that can result in the technical default of an otherwise sound issuer. Similarly, certain debt issues may be obligations not of a corporate parent but of a subsidiary, and their default therefore might not be considered a qualified *credit event* of the parent for the purposes of triggering a payout on a CDS. The International Swaps and Derivative Association (ISDA) has developed standardized documentation that defines six credit events that trigger payment on a CDS: bankruptcy, failure to pay,

obligation acceleration, obligation default, repudiation/moratorium, and restructuring. Readers should be aware that the terms are geographic-specific, with only bankruptcy and failure to pay (and, for some issues, restructuring) being recognized events for many contracts on North American entities. Obtaining legal opinion to review the terms of any OTC derivative, and credit derivatives in particular, is highly recommended.

7.9 HEDGING PREPAYMENT, REDEMPTION, AND OTHER HUMAN BEHAVIOR RISKS

In our discussion of operational risk in Chapter 2, we referred to the risk of losses that result from the behavior of employees and other individuals. These cases encompass fraud, error, and negligence. *Human behavior risk* refers to the risk of loss that arises from the actions taken by stakeholders, usually customers or investors, exercising their freedom of choice in the conduct of their affairs. Examples of such actions include the prepayment by a bank's customers of their mortgages and the early redemption or early termination of a customer relationship.

The effects on a lender of a customer's prepaying a loan obligation are twofold. We described in Chapters 2 and 3 how the ability of a mortgage borrower to exercise his option to terminate, or prepay, his mortgage contract exposes the mortgagee to reinvestment risk. In addition, if a financial institution had, in anticipation of the receipt of a series of cashflows, entered into a derivative contract to alter their characteristics (e.g., by swapping a series of floating rate cashflows into fixed, or vice versa) then any prepayment that resulted in the diminution or alteration of those cash receipts could result in basis risk due to a hedge mismatch. Similarly, a customer choosing to reduce the amount of his investment in a mutual fund would expose to basis risk a fund manager who had employed hedges to minimize the volatility of the stream of fund management fees.

Unfortunately, while individuals can be capricious and their behavior difficult to predict, human behavior risks, unlike market and credit risks, can be hedged only indirectly. A common approach used when attempting to hedge these risks is to understand the factors that drive behavior and to construct proxy hedges around these factors. For example, we have seen how mortgage prepayments are driven in large part by the prevailing level of interest rates, which determine the borrower's *refinancing incentive*, increasing as the market value of the loan exceeds the outstanding loan balance. Seasonality and seasoning—how long the loan has been in existence—also play a role in the refinancing decision, as does *burnout*, that is, whether the borrower has previously been exposed to rates that make refinancing

attractive. A bank seeking to hedge prepayment risk would use a prepayment model to estimate the change in prepayment activity for a pool of loans under a variety of future interest rate environments and the changes in portfolio cashflows and duration that would result. Analysts could then determine the sensitivity of the portfolio weighted average life and duration, for example, to a given change in future interest rate assumptions. Finally, interest rate derivatives could be used that would provide a payout in the event of an interest rate change that would be expected to cause a spike in prepayment activity. It should be apparent that sizing an appropriate prepayment hedge requires judgment, is fraught with model risk, and introduces basis risk. A financial institution considering such a hedge would have to decide if it is comfortable hedging prepayment risk and accepting the P/L volatility than such a hedge may introduce.

Redemption risk is faced by assets backed by investments, such as mutual funds and hedge funds, and bank deposits: An increase in redemptions (i.e., withdrawals) erodes the asset base. Mutual fund outflows are typically triggered by declining markets and can be hedged by taking a bearish market position. Like prepayments, bank deposit withdrawals are linked to interest rates and can be hedged with interest rate derivatives. Just as with prepayment risk, however, attempting to hedge redemption risk can introduce significant basis risk.

While prepayment and redemption risk can be estimated, it can be very difficult to forecast many other types of human behavior, with the only certainty being that the future is uncertain. Human behavior hedges cannot be exact and, rather than attempting to precisely forecast the future actions of stakeholders, firms should seek to identify all possible outcomes, however unlikely. Once the range of outcomes is identified as broadly as possible, estimates of the likelihood of each outcome can be made, as well as the impact of a given outcome. A common approach is to categorize vulnerabilities by estimating the impact of each identified potential outcome using a severity scale, for example, HS (high severity), MS (medium severity), and LS (low severity). A firm may choose not to hedge those outcomes considered to be highly probable but that are considered LS events, while it may decide that the survival of the firm requires that HS events, including those considered to be only remotely possible, must be hedged.

The key to successfully hedging human behavior risks is identifying those forces that influence individual decision making, and constructing a hedge around those factors. An airline, for example, could recognize that it is exposed to the risk that the crash of an airplane, even one belonging to a competitor, could lead vacationers to cancel their travel plans. To hedge its risk of loss from a reduction in passenger volumes, the airline might seek to purchase an OTC derivative or some other form of insurance to

protect against such a scenario. A related risk, one which it might not be possible to hedge, is a decline in travel in response to a terrorist act such as the September 11 attack in the United States. Because many insurance companies are unwilling to provide terrorism insurance, and there may be no counterparty willing to assume such a risk, a government may agree to insure such otherwise uninsurable risks. If not, it is a risk that the airline will have to bear.

7.10 EXECUTION

Once the decision to hedge has been made, and an appropriate hedge has been identified, it must then be executed, with the complexity of the execution process depending on the sophistication of the hedge and the nature of the instruments being used. The idiosyncrasies of individual instruments will be described in a subsequent volume, but common to the implementation of all hedges are general principles that should be adhered to, and questions that should be addressed. First, if the derivative is an over-the-counter instrument, then the identity of the hedge provider, and the size of the proposed exposure, should be ascertained in order to facilitate the management of counterparty risk. The precise nature of the protection being provided should be considered next. For example, what is the tenor of the hedge? What provisions exist for extending the maturity, or rolling over, of the hedge? Is the hedge a linear or a nonlinear instrument, such as a swap or an option position? If there are multiple risks to be hedged, is the hedge designed to mitigate more than one of the identified exposures and, if so, does the hedge consist of a single bespoke product or does the solution involve the use of multiple simpler instruments? For example, is there currency exposure to be hedged? If so, is a simple foreign exchange (FX) swap appropriate? Does an interest rate exposure in another currency have to be hedged, using for instance a cross-currency swap? Will an initial exchange of principal or a payment be required, as with, for example, an *off-market swap*, one where the fixed leg, such as in an interest-rate swap, is materially different from prevailing rates, and so a mark-to-market payment required at inception?

To properly analyze the benefit afforded by a hedge, its economics must be clearly understood. A starting point for this analysis entails an examination of the bid-ask spread for any instruments incorporated in the hedge, although with all but the most basic of hedges, this may be difficult to accomplish, with dealers seeking to maintain price obscurity. Nevertheless, a manager should require the dealer to confirm what volatility assumptions are used in pricing complex hedges, especially those involving multiple derivative

positions. For example, a *zero-cost collar*, also known as a *risk reversal*, involves the purchase of one option and the simultaneous sale of another. Using different volatility assumptions for the two options can allow a dealer to extract a hidden profit from an unsuspecting customer. The hedger would always do well to understand how the counterparty makes (or intends to make) money on the deal.

In general, the size of a derivative position should be specified in terms of its risk sensitivities, rather than notional amounts. For instance, when attempting to hedge the interest rate risk of a bond portfolio, knowing the notional amount of the bonds is not sufficient to determine the portfolio's sensitivity to interest rate changes. A portfolio of 30-year zero-coupon bonds, for example, will have a very different risk profile to a portfolio of five-year premium bonds. When describing a bond portfolio's risk profile, it is therefore necessary to specify the duration, in addition to the notional, of the portfolio with the risk sensitivity being a combination of the two factors, often being expressed in terms of the dollar change for a one basis point change in rates, or DV01. Similarly, when entering into an interest-rate swap a risk sensitivity such as DV01 is the appropriate way to specify the desired size of the hedge. Particular care should be paid to the other relevant sensitivities such as option Greeks (introduced in section 4.4) that, as discussed above, may have unintended consequences. For instance, one needs to pay attention to variance swap and option gamma and be aware of time decay (theta) when deploying equity options, for example, in a hedging solution.

As discussed earlier in this chapter, certain instruments such as futures contracts are marked to market daily, with their price changes often resulting in margin calls. It is important to both anticipate margin requirements, in order to prevent liquidity or cashflow problems and also to understand what procedures are in place for the daily marking of the hedge instrument, including the conventions used for pricing on those days when the market is closed, for example on holidays. Settlement conventions, including the number of days between trade date and settlement date, should be clearly understood. Similarly, it is important to be aware of and to understand the relevant day count convention, particularly for bonds, with typical conventions including actual/actual, actual/365 (government bonds), actual/360 (e.g., commercial paper and other short-term obligations), or 30/360 (e.g., agencies, corporates, and mortgage-backed securities). Managers of portfolios hedging foreign exposures by trading in markets in a different time zone (e.g., a U.S. manager hedging the Asian equity exposure of a portfolio of mutual funds by using the NIKKEI 225 index) should be aware that there will be timing differences in the marking to market of his U.S.

and Japanese positions, resulting in potentially significant short-term P/L volatility.

When entering trades, a customer with several accounts with an institution must specify the specific account into which the hedge is to be booked. Prior to executing a trade, the hedger should prepare a *trade ticket*, which includes the trade date, notional size, index or security, currency, counterparty firm and trader name, and the name of the person placing the trade. By entering all the relevant details on the ticket, except the actual exercise price and trade time, prior to trading, the hedger can use it as a reference when executing the trade to help him correctly specify the terms of his order. Market convention is that all phone calls on a trading desk are recorded, so that in the event of a dispute between the hedger and the trader regarding details of a trade the conversation can be replayed.

Once the order has been executed by the trader, the hedger can record the actual trade level on the ticket and sign and time-stamp it. Similarly, the hedger or his designate should enter trade details as soon as is practicable on a *trade blotter*, which is a spreadsheet or other computer program that is used to keep track of all trades entered into by the hedger. Ideally, the hedger will use a risk, reporting system, that contains the details of his portfolio and its risk sensitivities, and that can report position exposures. This system should be refreshed on at least a daily basis to reflect changing prices and market factors and should incorporate the risk sensitivities of both the manager's portfolio and the hedges that he has deployed, allowing him to see his aggregate net position and risk limits.

APPENDIX **A**

Basics of Probability Theory

PROBABILITY AND RANDOM PHENOMENA

Probability theory is concerned with the study of *random phenomena*, that is, those lacking *deterministic regularity*, while possessing at least some degree of *statistical regularity*. This means that while observations of random phenomena cannot be predicted with certainty, the observed *frequency* of the outcomes can in some sense be foretold. As is often the case in phenomenology, the qualifiers “with certainty” and “in some sense” are subject to interpretation, and the perceived deterministic or random nature of a particular phenomenon can depend on the meaning assigned to these qualifiers. One could say, in fact, that since, with a pair of well-known exceptions, nothing in life is certain, all phenomena are to some degree random.

In order to illustrate these concepts, let us examine that most prototypical of random phenomena, the tossing of a coin. In most people’s mind a toss of a coin can have one of two *equally likely* outcomes: heads or tails. The expectation of equal likelihood derives from *intuition* and, while usually taken as obvious, relies on a number of implicit assumptions. Leaving aside the unlikely but not inconceivable possibility of a third outcome (e.g., the coin could roll into a crack and become lost or wedged in a vertical position), these include (1) the coin being *fair*, or *unbiased*; (2) the tosses not being mechanically identical; and (3) the trials being independent. In practice, one or more of these assumptions may not pertain, at least not exactly. Consider the first one. It relies, in particular, on the coin’s being symmetric about its geometric center and on this center’s coinciding with the coin’s center of gravity. This may not be the case for a specie manufactured from two or more distinct alloys: The distribution of the two metals (having different density) may not be perfectly symmetric. In addition, the difference in designs embossed on the two sides could lead to asymmetric aerodynamic interaction with the surrounding air, favoring landing on one

particular side. Finally, the coin may simply be worn or damaged in one particular area.

The expectation of equal likelihood of the two outcomes also relies on an assumption that in each trial the coin is thrown differently or, even if not, that turbulent interaction with the air will obliterate the memory of any regularity in the throwing process. If one created a mechanical apparatus that launched the same coin repeatedly in exactly the same manner in a vacuum, then the coin would always land on the same side, say heads up, and the process would become deterministic, that is, the probability of heads would be exactly unity and of tails—nil.

In an experiment performed with this apparatus, the successive trials are perfectly dependent: A head is always followed by a head and a tail by a tail, depending on the way in which the coin is loaded. One might conceive of an imperfect version of this machine in which a head is likely, but not guaranteed, to be followed by a head (e.g., some residual air might introduce a small amount of turbulence). Then the probability of the next trial would depend (be conditioned) on the outcome of one or, in general, several preceding ones. The statistical description of this process would necessitate the use of *conditional probability*.

The above discussion demonstrates that the concepts of randomness and probability are to a large degree a matter of assumptions: Under the assumption of an unbiased coin and varying throws, the probability of both tails and heads is one-half; allowing for asymmetry in the coin's shape or for controlled tosses absent turbulence, the intuited probabilities would be different. A related conclusion is that probability theory cannot predict the future, even probabilistically: One can only make projections of expected frequencies of realized future outcomes based on a set of assumptions. The projections will come true to the degree to which the assumptions are correct and continue to hold during the observation period.

Consider now a researcher preparing to conduct an experiment with a coin assumed to be perfectly symmetric, sent flying through turbulent air by irreproducible kicks, and landing on a surface free of cracks. Based on these assumptions the researcher would calculate the *a priori* probability of both heads and tails to be one-half. This does not mean that after a large number N of trials, a head event will occur exactly $N_{\text{heads}} = \frac{1}{2}N$ times— N may not even be even! As N increases, however, the researcher would expect the *frequency* of realized heads events, defined as $f_{\text{heads}} \equiv N_{\text{heads}}/N$, to approach the heads' *a priori* probability $p_{\text{heads}} = \frac{1}{2}$:

$$f_{\text{heads}} \equiv \frac{N_{\text{heads}}}{N} \rightarrow p_{\text{heads}} = \frac{1}{2} \quad \text{as } N \rightarrow \infty. \quad (\text{A.1})$$

The discipline that studies realized data sets and their relationship to *a priori* probabilities is statistics; we will review some of its elements in Appendix B. In particular, statistics is concerned with *inferred probabilistic quantities* (e.g., the probability of heads as inferred from the outcomes of N trials) and their *statistical significance*. In our coin-tossing example, even though the measured frequency may not equal one-half, if a statistical significance test shows that f_{heads} is *compatible* with p_{heads} , then the researcher's assumptions are reasonable (although not proven). If, by contrast, the deviation of f_{heads} from p_{heads} is *statistically significant*, then the assumptions are probably wrong, for instance, the coin might be biased.

We emphasize that probability makes a statement about *statistical properties* of trial outcomes, not about individual trials. Even if the coin is unbiased, the tosses are non-deterministic, and the trials are independent, that is, if we know the exact probabilities of heads and tails, namely, $p_{\text{heads}} = p_{\text{tails}} = \frac{1}{2}$, we have no way of knowing what the outcome of the next, the second, or the n th next, trial will be. Moreover, even if we have observed 10 consecutive heads in the preceding 10 trials, the probability that the next toss of a fair coin will deliver a head is unchanged at $p_{\text{heads}} = \frac{1}{2}$ (we assume the trials to be independent), notwithstanding the fact that the probability of observing 11 consecutive heads is quite small.

As with any application of mathematics to physical (in particular, financial) phenomena, when studying probability, one should distinguish between a mathematical formalism and its correspondence to the real world. No matter how beguilingly elaborate a formalism might be, its conclusions cannot be better than the soundness of the assumptions linking it to observations.

PROBABILITY DISTRIBUTION FUNCTION

Consider again the case of a fair coin, where each trial (the toss of the coin) can yield one of two equally likely outcomes, heads (H) and tails (T), and the trials are independent—that is, the outcome of a given trial is unaffected by the outcomes of any of the preceding ones. Let $\text{Prob}\{A\}$ denote the probability of outcome A . In the case of a fair coin, A can assume one of two values: $A = H$ and $A = T$. Since, under our assumptions, one and only one of the two outcomes will occur in each trial and since the two outcomes are mutually exclusive (it is always *either* heads *or* tails), the probabilities of heads and tails should by conventional—but not unique—definition of probability sum to unity:

$$\text{Prob}\{H\} + \text{Prob}\{T\} = 1. \quad (\text{A.2})$$

At the same time, since we assume the heads and the tails to be equally likely, the corresponding probabilities should be equal:

$$\text{Prob } \{H\} = \text{Prob } \{T\}. \quad (\text{A.3})$$

Conditions in Equations (A.2) and (A.3) lead to

$$\text{Prob } \{H\} = \text{Prob } \{T\} = \frac{1}{2}. \quad (\text{A.4})$$

These probability values are a *consequence of our assumptions*.

It is often convenient to deal not with descriptive outcomes (heads/tails, no precipitations/rain/hail/snow, law passed/law rejected, buy/hold/sell) but with numbers. In the case of the toss of a coin, we could define a *random variable* ξ such that $\xi = +1$ if a trial yields a head and $\xi = -1$ if it produces a tail. We can then assign a probability to each value of ξ in accordance with equation (A.4), that is, define a *probability distribution function* (PDF) $p(\cdot)$:

$$p(\xi) = \begin{cases} \frac{1}{2}, & \xi = +1 \\ \frac{1}{2}, & \xi = -1 \\ 0, & \xi \neq \pm 1 \end{cases} \quad (\text{fair coin}). \quad (\text{A.5})$$

Similarly, for a *fair die* we could define ξ to equal the number of dots etched on each facet, leading to the following PDF:

$$p(\xi) = \begin{cases} \frac{1}{6}, & \xi = 1, 2, 3, 4, 5, \text{ or } 6 \\ 0, & \xi \neq 1, 2, 3, 4, 5, \text{ or } 6 \end{cases} \quad (\text{fair die}). \quad (\text{A.6})$$

Note that the choice of the mapping from trial outcome to number is arbitrary. For instance, in the case of the coin, we could have chosen $\{H \rightarrow 1, T \rightarrow 0\}$, or $\{H \rightarrow 0.6533, T \rightarrow -0.345\}$, or even $\{H \rightarrow 0, T \rightarrow 0\}$. The last, degenerate, case, although formally legitimate, would usually be of little practical use. Many-to-one mapping, however, can be meaningful and useful in more complex situations. Consider, for instance, a game where you win if the roll of a die produces one or two dots and lose otherwise. Then you might define $\xi = +1$ if a roll yields one or two and $\xi = -1$ if the

outcome is different. The corresponding probability distribution function would be

$$p(\xi) = \begin{cases} \frac{1}{3}, & \xi = +1 & (\text{you win}) \\ \frac{2}{3}, & \xi = -1 & (\text{you lose}). \\ 0, & \xi \neq \pm 1 \end{cases} \quad (\text{A.7})$$

Discrete Probability Distribution Functions

In both the coin and the die examples, there is a finite number of possible outcomes. When the set of possible outcomes is finite or when it is infinite but *countable* (i.e., when it is possible to establish a one-to-one correspondence between the possible outcomes and the—infinite—natural number series 1, 2, 3, ...), then the corresponding random variable and probability distribution are called *discrete*. For instance, an analyst might estimate the chances of the passage of a law at 60%. This is an example of a discrete distribution with two possible outcomes: $\text{Prob}\{\text{Law Passes}\} = 0.6$, $\text{Prob}\{\text{Law Rejected}\} = 0.4$. If we defined $\xi = +1$ when the law passes and $\xi = -1$ when it is rejected, then the probability distribution would be $p(+1) = 0.6$, $p(-1) = 0.4$, and $p(\xi) = 0$ if $\xi \neq \pm 1$.

Another example would be an economic projection with the probability of a base scenario ($\xi = 0$) estimated at 50% and the probabilities of an optimistic ($\xi = +1$) and a pessimistic ($\xi = -1$) scenario each at 25%. In this case $p(0) = 0.5$, $p(-1) = p(+1) = 0.25$, and $p(\xi) = 0$ if $\xi \neq 0, \pm 1$. (Of course, taken literally such probability estimates would imply spurious accuracy. They are often used by economists as a convenient distillation of a continuous—and known very imprecisely—distribution of possible economic trajectories.)

A frequently used discrete distribution with a countable infinite set of possible outcomes is the Poisson one. As we describe later in this appendix, the Poisson distribution can be used to describe the number of big price jumps in a given time period.

By definition, probabilities are non-negative and sum to unity (this is the normalization condition):

$$p(\xi) \geq 0 \quad \text{for all } \xi, \quad \sum_{\xi} p(\xi) = 1 \quad (\text{A.8})$$

where the summation extends over all values of ξ for which $p(\xi) \neq 0$.

Continuous Probability Distribution Functions

The probability distribution function of a discrete random variable gives the *probability* of this variable's assuming specific exact values—a discrete PDF is *concentrated* on a finite or countable infinite number of points and equals zero everywhere else. Some important variables, such as stock prices, asset returns, GDP figures, or unemployment rates, can assume a *continuum* of values (at least for practical purposes, if not in the strict mathematical analysis sense—see definition of continuous factors in section 3.3). For a *continuous random variable*, the probability distribution function $p(x)$ is not the probability that the variable equals x but rather the *probability density* that determines the likelihood of finding this variable in the vicinity of x . Specifically, for a small Δx , the probability that the value of the variable will fall between x and $x + \Delta x$ is given approximately by $p(x)\Delta x$. The exact statement is that the probability that a continuous random variable assumes a value between x_1 and x_2 equals

$$\text{Prob}\{x_1 < x \leq x_2\} = \int_{x_1}^{x_2} dx \, p(x). \quad (\text{A.9})$$

By definition, a probability distribution function is everywhere non-negative and obeys a normalization condition:

$$p(x) \geq 0 \quad \text{for all } x, \quad \int_{-\infty}^{\infty} dx \, p(x) = 1; \quad (\text{A.10})$$

this is a continuous generalization of equation (A.8).

Cumulative Distribution Function

Cumulative distribution function (CDF) $F(x)$ equals probability that the random variable assumes a value less than or equal to x :

$$\begin{aligned} F(x) &= \sum_{\xi \leq x} p(\xi) \quad (\text{discrete}), \\ F(x) &= \int_{-\infty}^x dx' \, p(x') \quad (\text{continuous}). \end{aligned} \quad (\text{A.11})$$

This definition and properties of probability distribution functions, stated in equations (A.8) and (A.10), imply that $F(x)$ is monotonically non-decreasing and is subject to limiting conditions:

$$\begin{aligned} F(x_2) &\geq F(x_1) \quad \text{for any } x_2 > x_1, \\ F(-\infty) &= 0, \quad F(+\infty) = 1. \end{aligned} \quad (\text{A.12})$$

If the CDF is known, then the probability that a random variable x falls between x_1 and x_2 can be calculated as a difference, rather than an integral, as in equation (A.9):

$$\text{Prob}\{x_1 < x \leq x_2\} = F(x_2) - F(x_1). \quad (\text{A.13})$$

A sample continuous PDF and a corresponding CDF are shown in Figures 5.4 and 5.5 respectively.

EXPECTED VALUES

For any function of a random variable $g(\cdot)$ the *expected value* $\langle g \rangle$ of this function is defined as

$$\begin{aligned} \langle g \rangle &\equiv \sum_{\xi} p(\xi) g(\xi) \quad (\text{discrete}), \\ \langle g \rangle &\equiv \int_{-\infty}^{\infty} dx p(x) g(x) \quad (\text{continuous}). \end{aligned} \quad (\text{A.14})$$

Intuitively, this is a sum (or integral—“continuous sum”) of function $g(\cdot)$ ’s all possible values multiplied by the likelihood of their occurrence. In a sequence of N trials the measured average value of $g(\cdot)$ should approach $\langle g \rangle$ with increasing N :

$$\frac{1}{N} \sum_{i=1}^N g(x_i) \rightarrow \langle g \rangle \quad \text{as} \quad N \rightarrow \infty. \quad (\text{A.15})$$

The speed of convergence depends on the details of the distribution function $p(\cdot)$, as well as the shape of the function $g(\cdot)$.

MOMENTS OF PROBABILITY DISTRIBUTION

Simple Moments

Simple moments m_ν (sometimes called just “moments”) of a probability distribution are expected values for a special case of $g(x) = x^\nu$ where ν is (usually but non necessarily) a positive integer:

$$\begin{aligned} m_\nu &\equiv \langle \xi^\nu \rangle = \sum_{\xi} p(\xi) \xi^\nu && \text{(discrete),} \\ m_\nu &\equiv \langle x^\nu \rangle = \int_{-\infty}^{\infty} dx p(x) x^\nu && \text{(continuous),} \\ \nu &= 1, 2, 3, \dots \end{aligned} \tag{A.16}$$

The first simple moment, m_1 , is known as the *mean*, the *average*, or the *expected value* of a random variable. We will alternatively denote it as $\langle x \rangle$ or $\bar{x}$.

Central Moments

Central moments μ_ν , $\nu = 2, 3, 4, \dots$, are expected values of a random variable’s deviation from its mean:

$$\begin{aligned} \mu_\nu &\equiv \langle (\xi - \bar{\xi})^\nu \rangle = \sum_{\xi} p(\xi) (\xi - \bar{\xi})^\nu && \text{(discrete),} \\ \mu_\nu &\equiv \langle (x - \bar{x})^\nu \rangle = \int_{-\infty}^{\infty} dx p(x) (x - \bar{x})^\nu && \text{(continuous),} \\ \nu &= 2, 3, 4, \dots \end{aligned} \tag{A.17}$$

Cumulants

Cumulants, also known, especially in mathematical literature, as semi-invariants, are particular algebraic combinations of simple moments. Their special significance and widespread use derive from the fact that a ν th cumulant of a sum of independent random variables equals the sum of these random variables’ ν th cumulants—hence the alternative name. We will denote the ν th cumulant of a random variable x as $c_\nu(x)$ or, when there is no ambiguity, simply c_ν .

A cumulant c_ν can be expressed as a sum, one of whose terms is m_ν and the rest are products of lower-order simple moments, multiplied by an

integer coefficient (which can, in general, be positive, negative, or zero). The sum of orders of moment in each term equals ν . While the exact expression for a general c_ν in terms of simple moments is straightforward, it is quite cumbersome, and an interested reader should consult any standard text on probability, such as Feller (1971) or Shiryaev (1995). An important point to remember is that the highest order simple moment involved in c_ν is of order ν .

Below we describe the first four cumulants, which are commonly used in financial economics and should suffice for most capital markets practitioners.

Mean The first cumulant, which we will often denote μ , is simply the *mean*: $c_1 \equiv \mu = m_1 = \langle x \rangle = \bar{x}$. It is the expected value of the random variable and is the number to which the statistical average of N observations, $N^{-1} \sum_{i=1}^N x_i$, converges as N increases.

Variance The second cumulant, commonly known as the *variance* and denoted σ^2 , coincides with the second central moment: $c_2 \equiv \sigma^2 = \mu_2$. It measures the degree of dispersion of the random variable about its mean. A more intuitive quantity is the square root of the variance, the *standard deviation* $\sigma \equiv \sqrt{c_2}$: Unlike the variance, the standard deviation can be directly compared to the mean.

Skewness The third cumulant, c_3 , which too coincides with a central moment, measures the degree of asymmetry of the distribution: A positive c_3 indicates that the distribution is skewed to the right, and a negative value, that it is skewed to the left. It is often useful to be able to express the third cumulant in terms of simple moments:

$$c_3 = \langle (x - \bar{x})^3 \rangle \equiv \mu_3 = m_3 - 3m_1m_2 + 2m_1^3. \quad (\text{A.18})$$

In practice, it is convenient to use the third cumulant scaled with the third power of standard deviation. This measure, γ , is known as (the coefficient of) *skewness*:

$$\gamma = \frac{c_3}{\sigma^3}. \quad (\text{A.19})$$

Kurtosis The fourth cumulant, c_4 , does not equal a central moment.¹ It measures the heaviness of the tails of the distribution. The cumulant c_4 is expressed in terms of central and simple moments as follows:

¹This is true for all c_ν with $\nu \geq 4$.

$$c_4 = \mu_4 - 3\mu_2^2 = m_4 - 3m_2^2 - 4m_1m_3 + 12m_1^2m_2 - 6m_1^4. \quad (\text{A.20})$$

The scaled version of the fourth cumulant is called *kurtosis*, κ :

$$\kappa = \frac{c_4}{\sigma^4} \equiv \frac{\mu_4}{\sigma^4} - 3. \quad (\text{A.21})$$

A positive kurtosis indicates that the distribution's tails are heavier than in a normal distribution—such distributions are called *leptokurtic*. A negative κ signifies lighter tails and a so-called *platykurtic* distribution.

Some readers may find this terminology confusing since *leptós* ($\lambda\epsilon\pi\tau\acute{o}\varsigma$) is Greek for “small,” “fine,” “light,” while *platús* ($\pi\lambda\alpha\tau\acute{\upsilon}\varsigma$)—“wide,” “broad,” “large.” As it happens, the “lightness” or “heaviness” in these names refers not to the tails but to the region around the peak, the *arch* of the distribution: *Kurtós* ($\kappa\upsilon\rho\tau\acute{o}\varsigma$) means “curved,” “arched,” “convex.” Since probability is “conserved” by virtue of normalization, equation (A.10), a heavier tail implies a lighter arch (hence leptokurtic) and a lighter, faster decaying tail necessitates more weight in the peak region (the distribution is, therefore, platykurtic).

A defining property of the normal distribution is that its skewness, kurtosis, and, generally, all the cumulants of order higher than the second vanish:

$$c_\nu = 0, \quad \nu = 3, 4, \dots \quad (\text{normal distribution}). \quad (\text{A.22})$$

The normal distribution is defined entirely by the mean and the variance.

Empirical Estimation of Higher Moments

It is clear from the definitions, equations (A.16) and (A.17), that a statistical estimate of moments, be they simple, central, or cumulants, becomes dependent on sampling progressively distant tails of the distribution as the order ν of the moment increases, as illustrated in Figure A.1, showing a sample distribution and the integrands for calculating the first four simple moments. An important practical consequence of this observation is that empirically calculated higher moments are typically less accurate than the lower ones since they are determined by the variable range where sampling is sparse.




FIGURE A.1 Simple moment integrands.

MEASURES OF CENTRAL TENDENCY

Measures of central tendency specify where data are concentrated.

Mean

We have already introduced the *mean*, defined by equation (A.16) with $\nu = 1$:

$$\mu \equiv \int_{-\infty}^{\infty} dx \, p(x) \, x \quad (\text{A.23})$$

(and a straightforward analog for a discrete distribution).

For a set of N observations of a random variable $\{x_i\}_{i=1}^N$ the statistical estimate of μ ,

$$\hat{\mu} \equiv \frac{1}{N} \sum_{i=1}^N x_i, \quad (\text{A.24})$$

is sensitive to *outliers*, that is, observations that are atypically large or small. An outlier may be a consequence of a one-off event, unlikely to recur, or simply an observational error. In addition, for distributions with sufficiently long tails (those decaying as $|x|^{-2}$ or more slowly) the mean does not exist² (or, rather, becomes infinite) and is, therefore, of little use.

Median

The *median* M of a random variable distribution is the value of the variable such that the chances of greater and smaller values are equal:

$$\int_{-\infty}^M dx \, p(x) = \int_M^{\infty} dx \, p(x) = \frac{1}{2}. \quad (\text{A.25})$$

In order to calculate a statistical estimate of the median from an observation set $\{x_i\}_{i=1}^N$ one should arrange x_i 's in ascending order. If N is odd, then the median estimate $\hat{M}$ is the middle element of the ordered set (the one occupying position number $(N+1)/2$). For an even set, $\hat{M}$ is the mean of the values, occupying two central positions, $N/2$ and $(N+2)/2$.

The median is more robust than the mean. For instance, in most practical situations it is virtually insensitive to outliers: The median of $\{-1\%, 2\%, 5\%\}$ is exactly the same as of $\{-13\%, 2\%, 5\%\}$, namely, 2% (the means are 2% and -2%, respectively).

Another useful property is that the median is well defined (and finite) for long-tailed distribution, including those with infinite mean.

Mode

The *mode* is the most probable value of a distribution. For a *unimodal* distribution, that is, one whose probability density function has a single maximum, the mode equals this maximum. When the density function has multiple maxima of equal height, then the mode is not unique.

Unlike the mean and the median, the mode can be used with non-numerical data. For instance, the distribution of surnames in the United States has a well-defined (and unique) mode, Smith. It is meaningless to talk of the “mean” or the “median” surname.

²If the distribution is symmetric, $p(-x) = p(x)$, then the integral defining the mean exists in the *principal value* sense, namely, as a limit of $\mu_b \equiv \int_{-b}^b dx \, p(x) x$ as $b \rightarrow \infty$. Since $\mu_b = 0$ for any b and, therefore, $\lim_{b \rightarrow \infty} \mu_b = 0$, one can say that in this very special sense the mean (as well as all the higher odd moments) of the distribution equals zero.

For a continuous distribution when all the three measures of central tendency exist and the probability density function is unimodal, either the mode, the median, and the mean coincide or the median is positioned between the mode and the mean.

SOME COMMONLY USED DISTRIBUTIONS

In this section, we list a number of frequently used probability distributions and briefly discuss their properties. An interested reader should consult any standard probability text, such as Feller (1971) and Shiryayev (1995), for more detail.

Normal Distribution

The *normal* or *Gaussian* distribution is the ubiquitous “bell curve” invoked to visualize the statistical properties of various quantities. Its defining property is that it is determined entirely by the first two cumulants, the mean μ and the variance σ^2 . All the higher cumulants equal zero, as stated in equation (A.22). The functional form of the normal probability density function is the exponential of minus the random variable’s scaled squared deviation from the mean, suitably normalized:

$$p(x) = \frac{1}{(2\pi\sigma^2)^{1/2}} e^{-(x-\mu)^2/2\sigma^2}. \quad (\text{A.26})$$

Since the normal distribution is symmetric, its median and mode coincide with the mean.

Log-Normal Distribution

The *log-normal* distribution governs a random variable whose logarithm is normally distributed. If natural logarithm $\ln x$ obeys a normal distribution with the mean μ and variance σ^2 , then x is governed by

$$p(x) = \frac{1}{(2\pi\sigma^2)^{1/2}x} e^{-(\ln x - \mu)^2/2\sigma^2} \quad \text{for } x > 0 \quad (\text{A.27})$$

and $p(x) = 0$ for $x \leq 0$.

The v th simple moment equals

$$m_v = e^{v\mu + \frac{1}{2}v^2\sigma^2}, \quad (\text{A.28})$$

whence the mean and the variance are given by

$$\begin{aligned}
c_1 &= e^{\mu + \frac{1}{2}\sigma^2} && \text{(mean),} \\
c_2 &= (e^{\sigma^2} - 1)e^{2\mu + \sigma^2} && \text{(variance).}
\end{aligned} \tag{A.29}$$

It is straightforward to calculate the skewness and the kurtosis using equations (A.18) and (A.20). The median is e^μ and the mode equals $e^{(\mu - \sigma^2)}$.

The log-normal distribution is frequently used to model the behavior of asset prices: A log-normal price, unlike a normal one, cannot become negative, and its logarithmic return is normal, making it amenable to analytical treatment.

Poisson Distribution

Consider discontinuous jumps in the price of an asset, of the kind we introduced in section 3.7 in the context of jump-diffusion processes. Suppose jumps are independent and let ξ be the number of jumps occurring during a given time interval $\mathcal{T}$. Let the expected value of ξ be $\langle \xi \rangle = \lambda$; this is the average number of jumps during interval $\mathcal{T}$. A natural question to ask then is, What is the probability that n discontinuous jumps will occur during $\mathcal{T}$?

The answer is given by the *Poisson* distribution:

$$\text{Prob}\{\xi = n\} \equiv p(n) = \frac{\lambda^n e^{-\lambda}}{n!}, \quad n = 0, 1, 2, \dots \tag{A.30}$$

where the factorial $n!$ is the product of all positive integers less than or equal to n : $n! \equiv 1 \cdot 2 \cdot \dots \cdot n$ (the factorial of zero is defined to be unity: $0! = 1$).

The value $p(0) = e^{-\lambda}$ is the probability that no jumps will occur, $p(1) = \lambda e^{-\lambda}$ is the probability that exactly one jump will occur, and the probability that *at least one* jump will occur is given by $1 - p(0) = 1 - e^{-\lambda}$.

A remarkable and useful property of the Poisson distribution is that all its cumulants are equal (and, consequently, equal to the expected value of ξ):

$$c_\nu = \lambda, \quad \nu = 1, 2, 3, \dots \tag{A.31}$$

Another way of putting this is that the Poisson distribution is defined by a single parameter, the *intensity* of the process, which can be interpreted as the expected value of the number of jumps.

MULTIVARIATE DISTRIBUTIONS

Let $\mathbf{x}$ be a collection, or *vector*, of K random variables:

$$\mathbf{x} \equiv \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_K \end{pmatrix}. \quad (\text{A.32})$$

One can view vector $\mathbf{x}$ itself as a K -dimensional random variable and define a probability density function $p(\cdot)$ similarly to the one-dimensional case: The probability of finding $\mathbf{x}$ in a small K -dimensional parallelepiped defined by a collection of intervals $[x_1, x_1 + \Delta x_1], \dots, [x_K, x_K + \Delta x_K]$ equals approximately $p(\mathbf{x})\Delta x_1 \Delta x_2 \dots \Delta x_K$. The precise statement, a multidimensional analogy of equation (A.9), is that for a region $\mathcal{V}$ in the K -dimensional space the probability that $\mathbf{x}$ assumes a value within this region equals the integral of $p(\cdot)$ over $\mathcal{V}$:

$$\text{Prob}\{\mathbf{x} \in \mathcal{V}\} = \int_{\mathcal{V}} dx_1 \dots dx_K p(\mathbf{x}) \equiv \int_{\mathcal{V}} d\mathbf{x} p(\mathbf{x}). \quad (\text{A.33})$$

If one considers $x_1, x_2, \dots, x_K$ as a collection of K (one-dimensional) random variables, then $p(\mathbf{x}) \equiv p(x_1, x_2, \dots, x_K)$ is their *joint probability density function*. Note that in general, the variables x_k are not independent. In fact, they are independent if and only if $p(\cdot)$ has the form

$$p(\mathbf{x}) \equiv p(x_1, x_2, \dots, x_K) = f_1(x_1) f_2(x_2) \dots f_K(x_K). \quad (\text{A.34})$$

Given $p(\mathbf{x})$, the one-dimensional probability density of each x_k (the so-called *marginal* probability density) can be obtained by “integrating out” all the other variables. For example, the marginal probability density of the first component of $\mathbf{x}$ is given by

$$p_1(x_1) = \int dx_2 \dots dx_K p(\mathbf{x}) \quad (\text{A.35})$$

where each integral extends from $-\infty$ to $+\infty$.

Covariances and Correlation Coefficients

In addition to the *diagonal* moments, defined by equations (A.16) and (A.17), along with the cumulants, for each of the K components of $\mathbf{x}$, multivariate distributions are also characterized by *cross-moments*, involving different components of $\mathbf{x}$. The most commonly used of these are *covariances*, measures of degree of x_k s' co-movement. The covariance of x_k and $x_{k'}$ is defined by

$$\text{cov}(x_k, x_{k'}) \equiv C_{kk'} = \int d\mathbf{x} \, p(\mathbf{x}) (x_k - \bar{x}_k) (x_{k'} - \bar{x}_{k'}). \quad (\text{A.36})$$

Note that when $k' = k$, the result is the variance σ_k^2 of x_k : Equation (A.36) defines a matrix with covariances in off-diagonal positions and variances along the main diagonal. This is the *variance-covariance matrix*:

$$\Sigma = \begin{pmatrix} \sigma_1^2 & C_{12} & \cdots & C_{1K} \\ C_{21} & \sigma_2^2 & & C_{2K} \\ \vdots & & \ddots & \vdots \\ C_{K1} & C_{K2} & \cdots & \sigma_K^2 \end{pmatrix}. \quad (\text{A.37})$$

The matrix Σ is symmetric: $C_{kk'} = C_{k'k}$, or $\Sigma^T = \Sigma$, where the superscript T indicates a transpose. It has $K(K+1)/2$ independent elements: half the off-diagonal ones, $(K^2 - K)/2$, plus K diagonals.

Covariances can be positive, negative, or zero (lack of co-movement). The magnitude of a non-zero covariance depends on the degree of variability of each of the variables, and, therefore, does not by itself convey the extent to which the moves of the variables are synchronized. The *correlation coefficient* corrects this deficiency by scaling the covariance with the random variables' standard deviations (assuming both σ_k and $\sigma_{k'}$ are non-zero):

$$\text{corr}(x_k, x_{k'}) \equiv \rho_{kk'} = \frac{\text{cov}(x_k, x_{k'})}{\sigma_k \sigma_{k'}} \quad (\text{A.38})$$

Unlike covariances, correlation coefficients can assume values between -1 (perfect anti-correlation) and $+1$ (perfect correlation), so the degree of co-movement is immediately clear.

Multivariate Normal Distribution

Multivariate normal is the multidimensional distribution, most commonly used in financial applications. For a K -dimensional random variable $\mathbf{x}$ with a vector of means $\boldsymbol{\mu}$ and variance-covariance matrix $\boldsymbol{\Sigma}$, the multivariate normal probability density function is

$$p(\mathbf{x}) = \frac{1}{[(2\pi)^K |\boldsymbol{\Sigma}|]^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}) \cdot \boldsymbol{\Sigma}^{-1} \cdot (\mathbf{x} - \boldsymbol{\mu}) \right] \quad (\text{A.39})$$

where the dot indicates matrix multiplication, $\boldsymbol{\Sigma}^{-1}$ denotes the inverse of $\boldsymbol{\Sigma}$, and $|\boldsymbol{\Sigma}|$ stands for its determinant.³

As in the case of one dimension, the normal distribution is completely determined by first- and second-order moments. All diagonal cumulants and cross-cumulants of orders higher than the second vanish.

One of the pragmatic reasons for the widespread use of multivariate normal distribution in financial applications is the relative ease of generating random vectors, obeying the distribution given by equation (A.39). The procedure involves representing the variance-covariance matrix as a product $\boldsymbol{\Sigma} = \mathbf{L} \cdot \mathbf{L}^T$ where $\mathbf{L}$ is a lower triangular matrix (i.e., a matrix whose elements above the main diagonal are zero). This can be accomplished by a procedure known as *Cholesky decomposition*. An interested reader can find a detailed description and practical algorithm in Press et al. (2002).

CENTRAL LIMIT THEOREM

The conceptual underpinning for the ubiquity of the normal distribution is the *central limit theorem* (CLT). The CLT states that, subject to certain constraints, the probability distribution of a sum of N independent random variables with finite second moments approaches a normal distribution as N increases. Specifically, if $\{x_i\}_{i=1}^N$ are independent random variables, $X_N \equiv \sum_{i=1}^N x_i$, $\bar{X}_N \equiv \sum_{i=1}^N \langle x_i \rangle$ is the expected value of X_N , and $\sigma_N^2 \equiv \sum_{i=1}^N \sigma_i^2$ its variance,⁴ then, as N increases, the distribution of

$$Y_N \equiv \frac{X_N - \bar{X}_N}{\sigma_N} \quad (\text{A.40})$$

approaches a normal distribution with zero mean and unit variance.

³A convenient refresher on matrix algebra is provided by Hamilton (1994), Appendix A.

⁴Note that the variance of a sum of random variables is equal to the sum of variances only when the variables are independent.

Speaking somewhat loosely, the CLT states that any quantity that can be expressed as a sum of reasonably independent finite-variance random variables is, to a commensurately reasonable approximation, normally distributed. In the context of capital markets this implies that the distribution of *logarithmic* returns should be approximately normal: The logarithmic return over a given time period is the sum of logarithmic returns over shorter periods; this is not true of arithmetic returns.

The applicability of the CLT to asset returns generally improves with increasing time period. For example, sub-second returns are more likely to be serially correlated (i.e., interdependent) than hourly ones and, therefore, a normal distribution is usually a better model for hourly (logarithmic) returns in liquid markets than, say, for returns over one second periods. This conclusion is corroborated by empirical observations as reported, for instance, in Mantegna and Stanley (2007) and Voit (2005).

The constraints in the formulation of the CLT place bounds on the “jolts” that individual variables can contribute to the sum. Note that we did not require that the variables be identically distributed. If we did, then there would be no need for additional constraints. Of course, in the context of capital markets it is precisely the prospect of atypical jolts that worries participants the most.

Rapidity of Convergence

The rapidity of convergence of the distribution of Y_N to the normal depends on the details of the distributions, governing the individual terms x_i . For a special case of *identically distributed* x_i s with a finite $g \equiv \langle |x_i|^3 \rangle$, the answer is given by the *Berry-Esséen theorem* (Shiryaev, 1995):

$$\sup_Y |F_N(Y) - \Phi(Y)| \leq \frac{A g^3}{\sigma^3 \sqrt{N}} \quad (\text{A.41})$$

where $F_N(\cdot)$ is the cumulative probability distribution (CDF) of Y_N , $\Phi(\cdot)$ is the CDF of the normal distribution with zero mean and unit variance, A is a numerical constant of order unity, and σ is the standard deviations of each x_i . The symbol “ $\sup_Y$ ” signifies *supremum*, which in most practical cases can be taken to mean “maximum.” Equation (A.41) says that the sum of identically distributed independent random variables *with a finite expected value of the cube of their magnitude* has a CDF whose deviation from the normal one decays as $N^{-1/2}$.

It is significant that the conditions for the Berry-Esséen theorem are considerably stricter than for the CLT: The variables must be identically distributed and $\langle |x_i|^3 \rangle$ must exist. While the additional conditions are likely to be satisfied, at least approximately, in many practical applications (although it is the exceptions that might be of the greatest interest, or concern), their existence is a manifestation of a general principle that it is easier to prove a statement in an unattainable theoretical limit (in this case $N \rightarrow \infty$) than to quantify the practically important pathway.

Relaxation of the Independence Condition

In practice, one has to deal with sums of random variables displaying various degrees of interdependence. This is especially true when dealing with financial time series, which are often prone to serial correlations.

It is possible to quantify the deviation of the sum's distribution from normal if the variables are "not too correlated" (Ma, 1985). The conditions on the random variables are mathematically expressed as bounds on all possible cross cumulants of all possible orders (to N), expressed as powers of N . If the conditions are satisfied, then Y_N of equation (A.40) has cumulants of order three and higher bounded by $N^{-1/2}$:

$$c_\nu(Y_N) \lesssim N^{-1/2} \quad (\nu = 3, 4, \dots) \quad (\text{A.42})$$

(for a normal distribution they would vanish exactly). If one is interested only in a few lower cumulants (the usual case), then one needs only check that the conditions are satisfied for cross cumulants up to the order of interest.

One should be aware, however, that the bounds are likely to be violated during abrupt market moves.

APPENDIX **B**

Elements of Statistics and Time Series Analysis

Statistics, “that gray area which is not quite a branch of mathematics—and just as surely not quite a branch of science” (Press et al., 2002), is primarily concerned with the description, study, and categorization of observed data sets and with relating these sets to theories regarding the data-generating processes. The statistical method typically consists in stating a hypothesis (e.g., “the correlation between two time series is zero”) and then assessing its plausibility given observations. Statistics can neither prove nor disprove a hypothesis but can, through a *statistical significance test*, produce a quantitative measure of the likelihood that the hypothesis is true.

A reader requiring the knowledge of statistics, its methods, and its specializations, such as time series analysis, should consult any of the numerous texts available. Here we only present a selection of results relevant to this book.

Asset Volatility

In the context of financial economics *asset volatility* usually means the standard deviation of this asset’s logarithmic returns. Let $\{S_i\}_{i=0}^N$ be a time series of an asset’s total return levels (in addition to price changes, they may incorporate the effects of distributions, such as dividends or coupon payments). Then a corresponding time series of logarithmic returns $\{r_i\}_{i=1}^N$ is defined by $r_i = \ln(S_i/S_{i-1})$, $i = 1, 2, \dots, N$.

In order to be able to compare directly the variability of returns, calculated with different frequencies (daily, weekly, monthly, quarterly, yearly, etc.), it is standard practice to *annualize* volatilities, that is, express them as if for each r_i the same rate of return pertained for a whole year, regardless of the actual frequency of observations. As an example, if an observed logarithmic return for a period of one calendar month were r_i , then the annualized

return, assuming the same rate of return for each month in a year, would be a sum of 12 identical r_i s, that is, $12r_i$. Since, as mentioned in Appendix A, the variance of a sum of independent random variables equals the sum of individual variables' variances, thus the *annualized variance* is defined as

$$\sigma^2 = \frac{\mathcal{A}}{N-1} \sum_{i=1}^N (r_i - \bar{r})^2, \quad \bar{r} \equiv \frac{1}{N} \sum_{i=1}^N r_i \quad (\text{B.1})$$

with $\mathcal{A}$ being the *annualization factor* equal to the number of times the return measurement period is contained in one year. Specifically, for yearly, quarterly, monthly, and weekly returns r_i , the annualization factor equals 1, 4, 12, and 52 respectively. For daily returns, $\mathcal{A}$ should be the number of business days in the year but is usually set to a fixed number between 250 and 260—when comparing volatilities of daily returns it is important to check whether the annualization factors are the same and make an adjustment if they are not.

As discussed after equation (5.1), if $\bar{r}$ is known *a priori* rather than being determined from the same data set as the variance, then one should change the denominator in the definition of σ^2 from $N-1$ to N .

To summarize, the volatility of an asset is the standard deviation of its annualized logarithmic returns, that is the square root of their annualized variance, equation (B.1):

$$\sigma = \left[\frac{\mathcal{A}}{N-1} \sum_{i=1}^N (r_i - \bar{r})^2 \right]^{1/2}. \quad (\text{B.2})$$

Linear Regressions

Linear regressions enable an analyst to quantify a perceived, suspected, or postulated linear relationship between some *dependent variable*, or *regressand*, y , and one or more *independent variables*, or *regressors*, or, especially in the econometric context, *factors*, $\{f_k\}_{k=1}^K$:

$$y = \alpha + \beta_1 f_1 + \beta_2 f_2 + \cdots + \beta_K f_K. \quad (\text{B.3})$$

This linear model assumes that a set of factors $\{f_k\}_{k=1}^K$ unambiguously determines the value of the dependent variable y via a collection of constant parameters, the *intercept* α and *slopes* β_k . In the language of section 3.3 (and somewhat deviating from its notation), f_k s are *continuous market* and

non-market factors. In fact, regressions can also involve discrete factors, and an interested reader should consult any standard statistics text.

The theory behind linear regressions imposes a number of preconditions on the observed time series (Hamilton, 1994). In practical applications these conditions are almost never exactly satisfied and, moreover, are usually difficult or impossible to ascertain. Linear regressions, nonetheless, are one of the most commonly used techniques in quantitative finance, and, when used judiciously and critically, can be a financial analyst's invaluable tool.

The role of statistics in linear regression is determining the values of parameters α and β_k from observed time series of dependent and independent variables. The most commonly used kind, the *ordinary least squares* (OLS) regression, determines the regression parameters by minimizing the sum of squared deviations from the model in equation (B.3). Specifically, let $y(t_i)$ and $\{f_k(t_i)\}_{k=1}^K$ with $i = 1, \dots, N$ be observed time series of, respectively, the dependent and independent variables. Then, given the values of parameters α and $\{\beta_k\}_{k=1}^K$, one can calculate the deviation $\epsilon(t_i)$ of the observed regressand from the model at each time t_i :

$$\epsilon(t_i) = y(t_i) - \alpha - \sum_{k=1}^K \beta_k f_k(t_i), \quad i = 1, \dots, N. \quad (\text{B.4})$$

The OLS method consists in finding α and $\{\beta_k\}_{k=1}^K$ that minimize the sum of squared deviations, that is, solving the following problem:

$$\min_{\alpha, \{\beta_k\}} E, \quad E \equiv \sum_{i=1}^N \epsilon^2(t_i). \quad (\text{B.5})$$

There exist numerous software applications for performing OLS regressions. In addition to calculating optimal parameters, they also output various significance statistics that quantify the degree to which the calculated α and β_k s are likely to reflect an actual relationship between the regressand and regressors, rather than occur by chance. One should be mindful, however, that these statistics, and the OLS method itself, presuppose a normal distribution of $\epsilon(t_i)$ s and may be incorrect or meaningless for data governed by long-tailed distributions. When long tails are present, *robust estimation techniques* might be more appropriate.

An interested reader can find an in-depth treatment of linear regressions in, for instance, Hamilton (1994).

Single Regressor

The single regressor, or univariate, case features prominently in the capital asset pricing model (CAPM), which states that the expected return on a given stock is linearly related to the expected return of the “market”:

$$r - r_f = \alpha + \beta(r_{\text{market}} - r_f). \quad (\text{B.6})$$

Here r is the expected return on the stock, r_f is the risk-free rate of return, and r_{market} is the expected market return. Linear regression is used to estimate α and β from realized stock and market returns.

Univariate regression is also important in the hedging context, especially in determining the size of the equivalent variation investment (EVI), as discussed in section 5.3.

The calculation of regression parameters for a univariate model, stated generally as

$$y(t_i) = \alpha + \beta f(t_i) + \epsilon(t_i) \quad (i = 1, 2, \dots, N), \quad (\text{B.7})$$

is particularly simple. The slope β can be expressed in terms of the observed correlation ρ between y and f and their respective standard deviations σ_y and σ_f :

$$\beta = \rho \frac{\sigma_y}{\sigma_f}, \quad (\text{B.8})$$

a result, already stated in equation (5.12). Here

$$\begin{aligned} \sigma_y^2 &= \frac{1}{N-1} \sum_{i=1}^N (y(t_i) - \bar{y})^2, & \sigma_f^2 &= \frac{1}{N-1} \sum_{i=1}^N (f(t_i) - \bar{f})^2, \\ \rho &= \frac{1}{(N-1)\sigma_y\sigma_f} \sum_{i=1}^N (y(t_i) - \bar{y})(f(t_i) - \bar{f}) \end{aligned} \quad (\text{B.9})$$

with the averages $\bar{y}$ and $\bar{f}$ calculated according to equation (A.24). Having thus estimated β , one can calculate a time series of residuals

$$e(t_i) \equiv y(t_i) - \beta f(t_i) \quad (\text{B.10})$$

and the slope α as their average:

$$\alpha = \langle e(t_i) \rangle \equiv \frac{1}{N} \sum_{i=1}^N e(t_i). \quad (\text{B.11})$$

Sometimes analysts need to constrain the intercept (usually by setting it to zero: $\alpha = 0$). In such cases equation (B.8) does not hold and one should solve for β by minimizing E in equation (B.5), as discussed below.

Multiple Regressors—Unconstrained This is a case of solving the general problem, stated in equations (B.4) and (B.5), with $K \geq 2$ and without any limitations on the possible values of the parameters α and $\{\beta_k\}$. There exists a closed-form solution that can be obtained by differentiating the sum E in equation (B.5) with respect to each parameter and equating the derivatives to zero. Solving the resulting system of linear algebraic equations for α and β_k s yields the desired parameters. This solution can be found in most standard statistics books and is provided by many numerical packages.

Multiple Regressors—Constrained In many situations the analyst needs to restrict the permissible values of regression parameters. The most commonly used set of parameter constraints forces the slopes to sum to unity and be non-negative:

$$\sum_{k=1}^K \beta_k = 1, \quad \beta_k \geq 0 \quad (k = 1, 2, \dots, K). \quad (\text{B.12})$$

These constraints enable one to treat β_k s as *portfolio weights*. For instance, if a manager wants to hedge his exposures with futures on a collection of indices, he might regress his portfolio's past returns on the returns of the chosen indices subject to constraints in equation (B.12) and then allocate exposures to these indices in proportion to β_k s.

Constrained multivariate regression is a nontrivial programming task and falls into a broader class of *non-linear constrained optimization* problems. An array of techniques for solving these problems has been developed, and they are available as part of many commercial computer applications and software libraries.

Coefficient of Determination (R^2) The *coefficient of determination*, also known as the R^2 or *R-squared*, is the fraction of the empirically observed regressand variance that is explained by the model:

$$R^2 = \frac{\sigma_{\hat{y}}^2}{\sigma_y^2} \quad (\text{B.13})$$

where

$$\sigma_{\hat{y}}^2 = \frac{1}{N-1} \sum_{i=1}^N (\hat{y}(t_i) - \langle \hat{y} \rangle)^2, \quad \sigma_y^2 = \frac{1}{N-1} \sum_{i=1}^N (y(t_i) - \langle y \rangle)^2, \quad (\text{B.14})$$

and the modeled value of the regressand is given by

$$\hat{y}(t_i) = \alpha + \sum_{k=1}^K \beta_k f_k(t_i). \quad (\text{B.15})$$

The averages in equation (B.14) are calculated as in equation (A.24).

In the case of unconstrained univariate regression of equation (B.7), the coefficient of determination can be shown to equal the square of the correlation coefficient between the dependent and independent variables:

$$R^2 = \rho_{yf}^2 \equiv [\text{corr}(y, f)]^2. \quad (\text{B.16})$$

The coefficient of determination can assume values between 0 and 1 inclusive with $R^2 = 0$ signifying that the model explains none of the regressand's variability and $R^2 = 1$ that all of it is explained.

The R^2 is an in-sample goodness-of-fit measure and as such does not say anything about the causal relationship between independent and dependent variables, the optimal number and composition of factors, or the model's predictive power. While a higher R^2 is generally desirable, one can usually increase the R^2 of a regression by simply adding new regressors. Such artificially bloated models tend to suffer from instability and multicollinearity, whose dangers we described in section 3.2, and mindless addition of new factors in order to boost the R^2 is likely to lead to an inferior model.

A number of modified versions of the coefficient of determination have been developed, which address some of these problems (e.g., by penalizing the addition of new factors).

Correlations

For two time series, $\{x_i\}_{i=1}^N$ and $\{y_i\}_{i=1}^N$, the statistical estimate of correlation coefficient, defined in equation (A.38), is given by

$$\hat{\rho} = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\left[\sum_{i=1}^N (x_i - \bar{x})^2\right]^{1/2} \left[\sum_{i=1}^N (y_i - \bar{y})^2\right]^{1/2}} \quad (\text{B.17})$$

where the averages $\bar{x}$ and $\bar{y}$ are calculated according to equation (A.24). Alternatively,

$$\hat{\rho} = \frac{1}{(N-1)\sigma_x\sigma_y} \sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y}) \quad (\text{B.18})$$

with

$$\sigma_x^2 = \frac{1}{(N-1)} \sum_{i=1}^N (x_i - \bar{x})^2 \quad (\text{B.19})$$

and similarly for σ_y^2 . The value of $\hat{\rho}$ lies between -1 and $+1$ inclusive.

An analyst may (and should) ask himself to what degree the estimate, produced by equation (B.17), reflects the actual correlation of random variables x and y and is not a statistical fluke, especially when the number of observations N is small.

Intuitively, if the value of $|\hat{\rho}|$ is close to unity, the estimate is more statistically significant, whereas if $|\hat{\rho}| \ll 1$, then variables x and y are likely to be uncorrelated and the measured value of $\hat{\rho}$ is probably statistical noise. Similarly, as N increases, so does the statistical significance of $\hat{\rho}$.

Statistics quantifies this intuition by devising a *statistical significance test*. In the case of correlations, the test begins with stating a *null hypothesis* that the correlation between the two time series is zero. The next step is to calculate the probability that, given the null hypothesis and the number of observations, a measured correlation magnitude would equal or exceed $|\hat{\rho}|$. If this probability, often called the *P-value*, is large, then the observation is compatible with the no-correlation assumption, and the data are insufficient to conclude that the time series are correlated with a coefficient of correlation equal to $\hat{\rho}$. If, however, the *P-value* is small, then a chance observation is unlikely to produce $\hat{\rho}$, and the statement that the correlation between x and y equals $\hat{\rho}$ is statistically significant.

Unfortunately, one needs to make further assumptions in order to calculate the probability distribution, governing $\hat{\rho}$. As is so often the case, the problem becomes analytically tractable if one assumes that x and y are normally distributed or, more accurately, that their joint distribution is *binormal*, that is, two-dimensional normal given by equation (A.39) with $K = 2$.

It turns out that under the assumptions of binormality and zero correlations the so-called *t-statistic*

$$t = \hat{\rho} \left[\frac{N-2}{1-\hat{\rho}^2} \right]^{1/2} \quad (\text{B.20})$$

obeys a particular distribution, known as *two-sided Student's t-distribution* with $N-2$ degrees of freedom (Press et al., 2002). The practical recipe is that in order to determine the statistical significance of a numerically calculated correlation coefficient, one should evaluate the *t*-statistic of equation (B.20) and calculate the value of the cumulative two-sided Student's *t*-distribution for this *t*-statistic and $N-2$ degrees of freedom (this CDF is available in most numerical packages and spreadsheets). If the resulting probability, the *P*-value, is of order unity, then the observation is compatible with the null hypothesis of zero correlation and $\hat{\rho}$ is not statistically significant. If, on the other hand, $P \ll 1$, then the observed correlation is unlikely to have occurred by chance and has a high statistical significance.

Note that the statistical test does not prove or disprove any assumptions or observations but provides a quantitative measure of the plausibility of the statement that the correlation coefficient between x and y equals $\hat{\rho}$.

This example is illustrative of the general statistical method: (1) State a hypothesis—the null hypothesis—usually, but not necessarily, incompatible with the observation being tested; (2) find a statistic whose probability distribution can be calculated—inevitably, under additional assumptions, typically of normality of random variables or their errors; (3) calculate the probability P of the empirical observation under the null hypothesis and reject the null hypothesis if this probability is small. There is no preset boundary in the *P*-value that separates acceptance from rejection—an analyst should be guided by the specifics of the problem and his objectives.

References

- M. W. Baxter and A. J. O. Rennie. *Financial Calculus*. Cambridge University Press, Cambridge, UK, 2001.
- G. Bekaert, C. R. Harvey, and A. Ng. Market integration and contagion. *Journal of Business*, 78:39–69, 2005.
- O. V. Bychuk and B. O'Shaughnessy. Anomalous diffusion at liquid surfaces. *Physical Review Letters*, 74:1795–8, 1995.
- O. V. Bychuk, A. Kurinets, and R. Schadt. *Testing Risk Models: Predictions and Hedging Performance*. FACTSET PMW User Conference, Atlanta, GA, 2002.
- N. Calamos. *Convertible Arbitrage: Insights and Techniques for Successful Hedging*. John Wiley & Sons, Hoboken, NJ, 2003.
- J. Y. Campbell and L. M. Viceira. *Strategic Asset Allocation*. Oxford University Press, Oxford, UK, 2002.
- L. Chew. *Managing Derivative Risks*. John Wiley & Sons, London, UK, 1996.
- R. Courant. *Differential and Integral Calculus*, 2nd Edition. Interscience Publishers, New York, 1952.
- J. C. Cox, S. A. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7:229–63, 1979.
- A. de Servigny and O. Renault. *Measuring and Managing Credit Risk*. McGraw-Hill, New York, 2004.
- K. Dowd. *Measuring Market Risk*. John Wiley & Sons, Chichester, UK, 2002.
- M. El-Erian. *When Markets Collide*. McGraw-Hill, New York, 2008.
- E. Elton, M. Gruber, S. Brown, and W. Goetzmann. *Modern Portfolio Theory and Investment Analysis*, 6th Edition. John Wiley & Sons, Hoboken, NJ, 2006.
- F. J. Fabozzi, editor. *The Handbook of Mortgage-Backed Securities*, 6th Edition. McGraw-Hill, New York, 2005.
- F. J. Fabozzi, editor. *Interest Rates, Term Structure, and Valuation Modeling*. John Wiley & Sons, Hoboken, NJ, 2002.
- F. J. Fabozzi and G. Fong. *Advanced Fixed Income Portfolio Management*. McGraw-Hill, New York, 1994.
- FDIC. *History of the Eighties—Lessons for the Future*. Federal Deposit Insurance Corporation, New York, 1997.
- W. Feller. *An Introduction to Probability Theory and Its Applications*, 2nd Edition. John Wiley & Sons, New York, 1971.
- P. Glasserman, P. Heidelberger, and P. Shahabuddin. Efficient Monte Carlo methods for value-at-risk. In J. Pickford and C. Alexander, editors, *Mastering Risk*, volume 2, pages 7–20. Financial Times—Prentice Hall, London, 2001.

- B. Graham. *The Intelligent Investor*. Harper & Row, New York, 1986.
- B. Graham and D. Dodd. *Security Analysis*. McGraw-Hill, New York, 2002.
- K. L. Grant. *Trading Risk*. John Wiley & Sons, Hoboken, NJ, 2004.
- J. D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1994.
- J. C. Hull. *Options, Futures, and Other Derivatives*, 7th Edition. Prentice-Hall, Inc., Upper Saddle River, NJ, 2008.
- B. I. Jacobs. *Capital Ideas and Market Realities*. Blackwell, Oxford, UK, 1999.
- P. Jorion. *Financial Risk Manager Handbook*. John Wiley & Sons, Hoboken, NJ, 2007.
- P. Jorion. Risk management lessons from Long-Term Capital Management. *European Financial Management*, 6:277–300, 2000.
- M. S. Joshi. *The Concepts and Practice of Mathematical Finance*. Cambridge University Press, Cambridge, UK, 2004.
- R. Lowenstein. *When Genius Failed*. Random House, New York, 2000.
- S.-K. Ma. *Statistical Mechanics*. World Scientific Publishing Co. Pte Ltd, Singapore, 1985.
- C. Mackay. *Extraordinary Popular Delusions and the Madness of Crowds*. Dover Publications, New York, 2003.
- B. B. Mandelbrot. The variation of certain speculative prices. *Journal of Business*, 36:394–419, 1963.
- R. N. Mantegna and H. E. Stanley. *An Introduction to Econophysics*. Cambridge University Press, Cambridge, UK, 2007.
- R. C. Merton. *Continuous-Time Finance*. Blackwell Publishers, Cambridge, MA & Oxford, UK, 1999.
- B. Øksendal. *Stochastic Differential Equations*, 5th Edition. Springer-Verlag, Berlin Heidelberg, 2000.
- J. Pickford and C. Alexander, editors. *Mastering Risk*. Financial Times—Prentice Hall, London, 2001.
- W. Preinitz. *A Fast Track to Structured Finance Modeling, Monitoring, and Valuation*. John Wiley & Sons, Hoboken, NJ, 2009.
- W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C++*. 2nd Edition, Cambridge University Press, Cambridge, UK, 2002.
- R. M. Rogers. *Handbook of Key Economic Indicators*. Irwin, New York, 1994.
- W. F. Sharpe. Asset allocation: Management style and performance measurement. *Journal of Portfolio Management*, 18:7–19, 1992.
- W. F. Sharpe. The Sharpe ratio. *Journal of Portfolio Management*, 21:49–58, 1994.
- A. N. Shiryaev. *Probability*. Springer-Verlag, New York, 1995.
- B. Solnik and D. McLeavey. *International Investments*. Pearson Addison Wesley, New York, 2003.
- F. A. Sortino and L. N. Price. Performance measurement in a downside risk framework. *Journal of Investing*, 3:59–64, 1994.
- S. Stapf, R. Kimmich, and R.-O. Seitter. Proton and deuteron field-cycling NMR relaxometry of liquids in porous glasses: Evidence for Lévy-walk statistics. *Physical Review Letters*, 75:2855–8, 1995.

- N. Taleb. *Dynamic Hedging*. John Wiley & Sons, New York, 1997.
- J. M. Tavakoli. *Collateralized Debt Obligations & Structured Finance*. John Wiley & Sons, Hoboken, NJ, 2003.
- B. Tuckman. *Fixed Income Securities*, 2nd Edition. John Wiley & Sons, Hoboken, NJ, 2002.
- P. Villarreal. Social security and medicare projections: 2009. *National Center For Policy Analysis*, 2009.
- J. Voit. *The Statistical Mechanics of Financial Market*, 3rd Edition. Springer-Verlag, Berlin Heidelberg, 2005.
- R. Yamarone. *Trader's Guide to Key Economic Indicators*. Bloomberg Press, Princeton, NJ, 2004.

Glossary

12b-1 Fees Back-end load—that is, deferred commission—fees on mutual funds, up to 1% per year for up to eight years. These fees, permitted under the SEC Rule 12b-1, are calculated on the net asset value (NAV) of the fund and are therefore subject to market risk. As of this writing, there is a pending proposal by the SEC to rescind Rule 12b-1 and replace it with Rule 12b-2, which would significantly alter the terms and structure of deferred commission fees.

ABS See *Asset-Backed Security*.

Absolute Delta The dollar sensitivity of the price of a derivative to a dollar change in the price of the underlying. See *Relative Delta*.

ADR See *American Depositary Receipt*.

Alpha In the capital asset pricing model, the portion of an asset's return uncorrelated with the return on the broad market and, therefore, diversifiable. In factor models, the portion of an asset's return uncorrelated with factor returns. In portfolio management, the portion of a portfolio's return uncorrelated to the portfolio's benchmark. Sometimes used loosely to mean the portion of a portfolio's return *in excess* of the benchmark return.

American Depositary Receipt (ADR) Certificate issued by an American bank and tradable on a U.S. exchange, representing an ownership interest in one or more shares of a foreign company.

American Option An option that can be exercised on any date during its life, in contrast to a European option, which can only be exercised on its expiration date.

Amortizing Bond A bond that pays principal on a periodic basis, in addition to regularly scheduled interest. See *Bullet Bond*.

Amortizing Swap See *Balance Guarantee Swap*.

Arbitrage A risk-free strategy that results in a profit, such as the simultaneous purchase and sale of a security at different prices. It depends on assets being mispriced.

Arbitrage Spread The spread between two identical, or highly similar, grades of oil, in two different locations.

Asian Option An option whose payoff is a function of the average price of the underlying over some specified period of time.

Ask Price The price demanded by a seller. See *Bid Price*.

Asset-Backed Security A bond whose interest and principal payments derive from cash flows generated by the security's collateral such as auto loans, credit cards, home equity loans, etc. Typically, an ABS is arranged by a corporate issuer such as a bank, but is issued by a special purpose entity, and the credit risk

is a function solely of the collateral, unless the bond is insured. In contrast, a corporate bond is an obligation of the corporate issuer.

At-the-Money Option An option whose exercise (strike) price equals the current price of the underlying asset.

Available-for-Sale Security A security that is neither a trading security nor a hold-to-maturity security and is held with the expectation of being sold in either the short or the long term. It is reported at fair value, with changes in valuation recognized in equity rather than net income.

Back Office The back office in a financial organization handles trade settlement and reconciliation, compliance, and financial and regulatory reporting. See *Front Office* and *Middle Office*.

Back-Testing Testing a quantitative model, such as a risk or trading model, using historical data to examine the model's explanatory or predictive power.

Backwardation A situation in futures markets when the price curve is inverted, with prices being lower in distant delivery months.

Balance Guarantee Swap An amortizing swap, one in which the notional amount declines over time to match the principal balance of some reference asset. Such a swap could be used to hedge a mortgage portfolio, for example, in which the principal balance declines over time due to both the scheduled retirement of principal and prepayments.

Barrier Option A path-dependent option whose payout depends on the price of the underlying asset reaching or passing a defined threshold. There are two basic types: A knock-in option becomes active once the specified barrier, or trigger, price is reached, at which time the knock-in functions like a plain vanilla option. A knock-out option ceases to exist once the barrier price is reached by the underlying asset. See *Up and In*, *Up and Out*, *Down and In*, and *Down and Out*.

Basis The difference between the spot and futures prices of a commodity. The basis will almost certainly change over time.

Basis Point One hundredth of one percent. Often abbreviated to *b.p.* or *bp* and pronounced "bip" (plural: *bips*, pronounced "bips").

Basis Risk The risk of loss with a hedged position arising from a difference in risk sensitivities between the position to be hedged and the hedging instrument.

Basis Swap An interest rate swap where the two legs relate to floating reference rates.

Basket Swap A swap where one leg consists of a portfolio of assets or reference entities. For example, a credit default swap on a basket of bonds.

Behavioral Finance The study of psychological aspects of the investment process. Some students of financial markets have proposed that investors may behave irrationally (or apparently irrationally), resulting in market price anomalies.

Benchmark A diversified portfolio or index that forms a reference against which the performance of a portfolio, and its manager, may be measured. Also referred to as *bogey*.

Benchmark Risk The risk of loss to a portfolio relative to a benchmark, arising from the fact that the benchmark may not be fully investable or because a manager will only be able to invest in the component securities at prices prevailing after the index composition was made public.

- Bermudan Option** An option that can be exercised only on specified dates over its lifetime.
- Beta** A measure of the systematic risk of an asset or portfolio due to its covariance with the market, that is, the degree to which its returns move with the market. According to theories such as the Capital Asset Pricing Model (CAPM), the beta can be used to calculate the expected systematic return on an asset or portfolio.
- Bid-Ask Spread** The difference between the bid and the ask prices.
- Bid-Offer Spread** See *Bid-Ask Spread*.
- Bid Price** The price offered by a prospective purchaser. See *Ask Price*.
- Black Box** A model where only the inputs and outputs are available, with the implementation and model design being invisible to the user.
- Black-Scholes Option Pricing Model** A model for pricing plain vanilla European options that admits a closed-form solution.
- Bogey** See *Benchmark*.
- Bond Insurance** A form of insurance that guarantees the payment of interest and principal due on a bond. If the issuer is unable to make a payment, then the insurer is expected to satisfy the obligation. Bond insurance causes the credit rating of a bond to equal or exceed that of the insurer.
- Bootstrapping** A technique of calculating spot (zero-coupon) rates and forward rates from a par coupon curve.
- Breaking the Buck** An instance of the NAV of a money market fund falling below one dollar.
- Bullet Bond** A bond that pays interest at regular intervals and principal at maturity.
- Busted Convertible** A convertible bond trading at a price so far below its conversion price that the conversion feature has no value, so that the bond is priced like a plain-vanilla bond.
- Callable Bond** A bond that can be called, or redeemed early, at a specified price, at specified times, by its issuer. Because of the risk of early exercise, the yield on a callable bond—which is effectively short a call option—will usually be higher than the yield on a plain-vanilla bond of the same rating and maturity.
- Call Option** An option that conveys the right, but not the obligation, to purchase an asset at a specified price by or on a specified date.
- Cap** See *Interest Rate Cap*.
- Capital Asset Pricing Model** A theory stating that the expected return on an asset equals the return on the risk-free asset, plus the market risk premium multiplied by the asset's beta, a measure of the asset's co-movement with the market.
- Caplet** A constituent of an interest rate cap.
- CAPM** See *Capital Asset Pricing Model*.
- Cash Flow Waterfall** See *Waterfall*.
- CDO** See *Collateralized Debt Obligation*.
- CDS** See *Credit Default Swap*.
- Charm** The cross-sensitivity of price and time. It can be viewed as the sensitivity of delta to a change in time, or theta's sensitivity to a change in price.
- CLO** See *Collateralized Loan Obligation*.
- CMO** See *Collateralized Mortgage Obligation*.

Coefficient of Determination (R^2) An in-sample goodness-of-fit measure of a data model, namely, the fraction of the dependent variable's variance explained by the model. Can assume values between 0 (model explains none of the dependent variable's behavior) and 1 (model explains all of the observed variability). In case of unconstrained ordinary least squares, linear regression equals the square of the correlation coefficient between the observed and the modeled values. If, in addition, the regression is one-dimensional, then also equals the square of the correlation coefficient between the dependent and independent variables.

Collar See *Interest Rate Collar*.

Collateralized Debt Obligation An asset-backed security in which the collateral consists of a portfolio of bonds, each of which can have a different credit rating.

Collateralized Loan Obligation An asset-backed security in which the collateral consists of loans, usually corporate obligations.

Collateralized Mortgage Obligation An asset-backed security in which the collateral consists of a pool of mortgages, with interest and principal payments flowing through a waterfall to a series of constituent classes, or tranches, that can have different target maturities.

Color A risk sensitivity of nonlinear assets and portfolios, it is the rate of time decay of the gamma.

Commodity Risk The risk of loss to a business arising from an adverse move in the price of commodities.

Compliance Risk The risk of loss due to censure and penalties arising from violations of the law and from a failure to properly supervise employees.

Compound Option An option on an option, such as a call option that gives the right to purchase a put option.

Contagion Risk The risk that adverse market conditions for one market sector will spread to other instruments or markets.

Contango A situation in which futures prices are higher than spot prices and futures prices for more distant delivery dates are higher than for nearer delivery dates.

Contingent Claim An obligation that may exist in the future, such as the requirement by the seller to perform on an in-the-money call option, if exercised.

Continuous Compounding A method for calculating interest, based on infinitesimally short time periods.

Convertible Bond A bond that can be converted to equity under certain conditions, typically if the price of the equity increases above a certain threshold. Because conversion rights have economic value, such bonds are usually sold with coupons lower than those offered by plain-vanilla bonds, in return for the chance to earn a higher return if the bond converts.

Convexity A second-order sensitivity, it is a measure of the sensitivity of a bond or portfolio's duration to changes in yield. Positive convexity is a phenomenon whereby the price increase that would result from a decrease in yield is greater than the price decrease that would result from a yield increase of the same magnitude. Negative convexity describes the opposite phenomenon and is generally unattractive to bond investors. Callable bonds exhibit negative convexity in low interest rate environment.

- Correlation Coefficient** A measure of how two portfolios move together. It is equal to their covariance, scaled by the standard deviations of the two portfolios. It has a range of -1 , signifying that the portfolios move in opposite directions, to $+1$, signifying that they move perfectly together. A value of 0 means there is no co-movement.
- Correlation Risk** The risk of loss due to a change, particularly an unexpected movement, in the relationship between one or more instruments or positions.
- Cost of Carry** The cost of financing an asset, plus any associated storage costs, less any income derived from the asset. Unlike financial assets, physical assets such as oil can have significant storage costs, as well as being subject to evaporation and spoilage.
- Cost to Close** The cost of unwinding or closing out a position.
- Counterparty Limit** A maximum allowed exposure to a given counterparty, often set by a firm as a risk control measure.
- Counterparty Risk** The risk of loss due to a failure to perform on the part of a counterparty. Unlike default risk, which relates to a failure to repay a credit obligation such as a loan, counterparty risk relates to a failure to honor obligations on a contract, such as an in-the-money swap or option position.
- Country Risk** See *Political Risk*.
- Covariance** A measure of the degree to which two assets or portfolios move together.
- Covered Call** A short position in a call option combined with a long position in the underlying asset.
- Credit Default Spread** The price of default protection on an issuer, it is a measure of the issuer's default risk.
- Credit Default Swap** An instrument that provides the owner, in return for the periodic payment of a premium—the credit default spread—with the right to sell a bond for its full par value in the event of a default or other prespecified credit event by the issuer.
- Credit Derivative** A derivative instrument whose payoff depends on the credit-worthiness of an issuer. For example, a credit default swap, which provides a payout in the event of a bond default by the “name,” that is, the specified issuer.
- Credit Event** A default, or failure to perform, by an issuer on one of its obligations. The types of nonperformance, and the specific bonds or other obligations on which nonperformance may be considered an “event,” are usually clearly defined in a credit derivative contract. For example, the International Swaps and Derivative Association (ISDA) defines six credit events that trigger payment on a CDS: bankruptcy, failure to pay, obligation acceleration, obligation default, repudiation/moratorium, and restructuring. Terms can be geographic-specific, with only some events being applicable in certain jurisdictions.
- Credit Rating** A measure of an obligor's creditworthiness that can be used to derive the obligor's expected probability of default. A rating can relate to the expected ability of the issuer—an *issuer rating*—or a given obligation—an *issue-specific rating*—to perform.

- Credit Risk** The risk of loss due to the failure of a borrower or other obligor to honor its obligations. See *Spread Risk*.
- Credit Spread** The difference in yield between two securities of the same maturity, usually a corporate bond and a government bond, resulting from their different perceived levels of credit risk.
- Credit Spread Option** An option whose payoff is a function of the spread between the obligations of an issuer and a government bond of the same maturity, and that is typically used to protect against credit and spread risks.
- Cross-Currency Swap** A swap whereby interest payments in one currency are exchanged for those of another. At swap inception and maturity, there is an exchange of principal, leading to the potential for significant counterparty risk.
- Cross-Default Provision** A clause in a contract specifying that the default by an obligor on one debt obligation will trigger a default on another.
- Cross-Hedge** A hedging strategy whereby the market risk of a cash asset is hedged with a derivative on another, ideally closely correlated, asset. An alternative meaning is that a cross-hedge is a forward contract to sell one foreign currency and receive another currency. The hedger still has foreign currency exposure, merely replacing one exposure with another, in the expectation that one currency will outperform the other.
- Cumulant** An algebraic combination of probability distributions' simple moments that is invariant under the addition of independent random variables. Also called *semi-invariant*.
- Curtailement** The process whereby a mortgagor makes sporadic or regular payments of additional principal in order to shorten the final maturity of the mortgage.
- Deflation** A general decrease in the prevailing level of prices for goods and services.
- Delinquency Rate** The percentage of loans in a portfolio where the borrower has failed to keep the loan current, that is, has missed one or more payments.
- Delta** The sensitivity of the price of a derivative to a change in price of the underlying.
- Delta Decay** See *Charm*.
- Delta Hedging** A hedging strategy that seeks to make the value of a portfolio of assets insensitive to small changes in the price of the underlying by buying or selling futures or other derivatives. A portfolio of derivatives can be delta-hedged by buying or selling short underlying instruments.
- Delta Neutral** A delta-hedged portfolio, that is insensitive to small changes in the price of the asset.
- Delta Riding** A method of accumulating gain as part of the effective delta hedging strategy whereby a positive-gamma delta-neutral portfolio is allowed to build material delta as a consequence of market moves before an adjustment to delta-neutral. Also called *gamma climbing*.
- Derivative** A financial instrument whose value derives from the change in price or other financial characteristic of some other, underlying asset or instrument.
- Derivatives Risk** The risk of loss arising from an inappropriate or poorly implemented derivatives strategy.
- DgammaDvol** See *Zomma*.

- Discounted Cash Flow (DCF)** A method of valuing a security by discounting to the present the projected future cash flows due to the holder of the security.
- Discount Rate** The interest rate at which the Federal Reserve will lend to member banks.
- Diversification** The process of incorporating a variety of negatively, or less than perfectly, correlated securities in an investment portfolio in order to increase its risk-adjusted return by reducing its risk.
- Dollar Value of a Basis Point** See *DV01*.
- Dotcom Bubble** A speculative bubble in technology stocks characterized by a rapid increase in stock prices between 1996 and 2000, followed by a calamitous decline over the next three years, with the Nasdaq 100 index of stocks dropping over 80% between March 31, 2000, and September 30, 2002.
- Down and In** A type of knock-in barrier option, one which is activated by the underlying asset price declining through the barrier price.
- Down and Out** A type of knock-out barrier option, one which ceases to exist after the underlying asset price declines through the barrier price.
- Drawdown** The maximum loss realized since the previous peak portfolio value. It is sample dependent, unlike the mean and variance.
- Duration** Can refer to a measure of the weighted average time to the receipt of cash flows for a bond, or to the sensitivity of the value of a set of cash flows such as those of a bond or swap to a change in interest rates. See *Macaulay Duration*, *Modified Duration*, and *Effective Duration*. “Duration” usually is taken to refer to modified duration.
- DV01** The dollar change in value of a portfolio or asset in response to a one basis point move—usually a decline—in interest rates. The usual North American name for *PV01*.
- DvegaDtime** A second order sensitivity, it is the measure of the change in vega with the passage of time, or the change in theta with a change in volatility.
- DvolgaDvol** See *Ultima*.
- Dynamic Hedging** A strategy to hedge a portfolio of derivatives by periodically rebalancing a position in the underlying as prices or other conditions change.
- Earnings Risk** The potential for a company’s earnings to be lower than anticipated.
- Earnings-at-Risk** An estimate of the reduction in earnings, or net income, for an enterprise that would result from an adverse move in interest rates.
- Economic Indicators** Statistics about the level of economic activity. Some may have predictive power and are referred to as *leading* indicators. Others relate to historical behavior and are called *trailing*, or *lagging*, indicators. Finally, indicators that provide information about the current state of the economy are known as *coincidental* indicators.
- Effective Delta** The change in a derivative’s or portfolio’s value in response to a *finite* move in the price of an underlying security. Should incorporate the effects of likely concomitant co-movement of other factors (e.g., volatility). Unlike traditional delta, the effective delta is a function of not only the current level of the price and other factors but also of the *size* of the price move (in general, of two distinct sizes for up and down moves). Also called *modified* or *shadow* delta.

Effective Duration A measure of the sensitivity of the price of a bond or other fixed-income instrument to a change in interest rates. Unlike *modified duration* (*q.v.*), the effective duration takes into account the fact that a change in interest rates can alter the instrument's cash flows. With a mortgage-backed security, for example, a change in interest rates is likely to lead to a change in prepayment behavior.

Elasticity See *Omega*.

Equity Risk The risk of loss due to an adverse movement in the price of an equity security.

Equity Swap See *Total Return Swap* and *Price Swap*.

Equivalent Variation Investment (EVI) A position in a market factor instrument that can be used to hedge the market risk of an asset. The position is sized such that the variations in the position's value for small moves in the market factor equal the asset's value changes attributable to these moves.

Error Function Cumulative probability distribution function $\Phi(x)$ for a standardized normal distribution. That is, $\Phi(x)$ is the probability that a random variable, governed by a normal distribution with zero mean and unit variance, will be less than x : $\Phi(x) \equiv (2\pi)^{-1/2} \int_{-\infty}^x d\xi \exp(-\xi^2/2)$. The importance and ubiquitousness of the error function in quantitative finance stems from the fact that it arises naturally and features prominently in exact and approximate derivative pricing formulae when the price of the underlying is assumed to obey the log-normal distribution.

European Option An option that can only be exercised on its expiration date, in contrast to American and Bermudan options.

Exchange Risk The risk of loss due to a failure by an organized exchange, whose members are jointly and severally liable for the exchange's obligations.

Exercise Price The price at which an option allows the underlying to be purchased or sold. Also known as strike price.

Exotic Option An option contract with nonstandard features.

Expected Return According to a model such as the CAPM, it is the return that investors in a security would expect based on the security's risk relative to the market. It can also refer to the sum of possible returns for an instrument weighted by their probabilities.

Expected Tail Loss (ETL) The mean loss expected in the tail of a distribution. Whereas the value at risk indicates the maximum loss that is expected to occur, at a given confidence interval, the ETL indicates the magnitude of the average loss that is expected to occur with a frequency of one minus the confidence interval. For example, if the confidence interval is 95%, the ETL is the mean loss expected to occur in the 5% of cases where the loss is expected to exceed the value at risk.

Expected Value The sum of the possible outcomes of a variable multiplied by their probabilities.

Exposure Hedge Ratio The percentage of a portfolio that is to be hedged.

Exposure Limit See *Counterparty Limit*.

Extension Risk The risk that the duration of an obligation will lengthen. For example, the duration of a mortgage-backed security may increase due to the rate of mortgage prepayments being less than anticipated.

- Fair Value** The price at which an asset or liability could be exchanged between willing participants, other than in a liquidation. In the absence of an active market, the fair value may only be an estimate, leading to the potential for disagreement over what value should be reported on the books of a firm.
- Federal Open Market Committee (FOMC)** A committee of the U.S. Federal Reserve System—the Fed—that oversees the open market operations by which the Fed controls monetary policy.
- Financial Accounting Standard (FAS)** One of a number of procedures promulgated by the Financial Accounting Standards Board (FASB), defining the accounting treatment of transactions and the valuation of financial assets and derivatives.
- Floating Rate** An interest rate that changes over time. Usually it is calculated as a spread over a reference rate such as the LIBOR rate for a particular tenor. For example, a rate of 25 basis points over three-month LIBOR might be quoted as “3ML+25.”
- Floor** See *Interest Rate Floor*.
- Foreclosure Rate** Percentage of mortgage loans in foreclosure, that is, where the mortgagee is in the process of seizing, or has already seized, a property due to a default by the mortgagor.
- Foreign Exchange Risk** The risk of loss due to an unfavorable move in currency exchange rates.
- Forward Contract** An over-the-counter contract that obligates the holder to buy or sell a specified quantity of a specified asset, at a specified price, on a specified date. See also *Futures Contract*.
- Forward Rate** The interest rate for a future time period that is implied by the prevailing spot rates.
- Front Office** Traditionally, the client-facing area of a financial organization. It generally refers to the sales and trading function and associated profit centers. See also *Back Office* and *Middle Office*.
- Funded Hedge** A hedge position consisting of instruments such as put options, where the premium has been paid and no mark-to-market payments will be required in the future.
- Futures Contract** A contract that obligates the holder to buy or sell a specified quantity of a specified asset, at a specified price, on a specified date. Unlike a forward contract, a futures contract has standardized terms, is listed on an exchange, with the counterparty being the exchange, and is marked to market daily.
- FX Risk** See *Foreign Exchange Risk*.
- Gamma** A second-order sensitivity, the gamma is a measure of the price sensitivity of delta, that is, the sensitivity of the sensitivity of the value of an asset, usually an option, to a change in price of the underlying.
- Gamma Climbing** See *Delta Riding*.
- Gamma Hedging** A hedging strategy in which the gamma, as well as the delta, is kept small. A gamma-hedged portfolio’s delta changes little with changes in the markets, thus necessitating less frequent delta-hedge rebalancing and incurring reduced transaction costs. For assets with a positive gamma, this strategy forgoes or reduces the opportunity for delta riding.

Gate Provision A restriction placed by an asset manager on the magnitude and timing of investor redemptions.

Greater Fool Theory A belief held by an investor in an overpriced security, usually a non-dividend-paying stock, that there will be a buyer, expecting an even greater price appreciation, to whom the investor will be able to sell the security at a profit. In a typical stock market bubble, an equity passes through a chain of ever greater fools until, as the bubble bursts, the last investor in the chain is left with an asset he cannot sell except at a great loss.

Greeks A collection of market risk sensitivities, so called because many of them are denoted with Greek letters. These sensitivities are commonly used to measure and manage the risk of nonlinear instruments, such as options and delta-hedged portfolios. The Greeks can be first-order, such as the delta, the vega, the up-silon, the rho, and the theta; second-order, such as the gamma, the charm, the DvegaDtime, the vanna, and the volga; and third-order, such as the color, the speed, the ultima, and the zomma.

Haircut A discount applied to the value of securities held as collateral, in order to protect the lender against an adverse price move. See *Repurchase Agreement*.

Headline Risk The risk of loss resulting from market reaction to unfavorable news about a company in the mass media.

Hedge Ratio The ratio of the size of a position in hedging instruments such as futures to the size of the cash position being hedged.

Historical Volatility Volatility estimated from a historical time series.

Hold-to-Maturity Security A debt security held by an entity that has the ability and the intention to hold it until maturity. It can be classified as a current asset or a long-term asset, depending on its maturity and is reported at amortized cost. Changes in the fair value of such a security are not reported. Note that equity securities, having no maturity, cannot be classified as hold to maturity.

Human Behavior Risk The risk of loss due to the behavior of individuals, such as a default by borrowers whose mortgages underlie a collateralized mortgage obligation or the loss of management fees for a mutual fund manager whose clients redeem their shares.

Idiosyncratic Risk See *Nonsystematic Risk*.

Implied Volatility Price volatility implied by a model such as the Black-Scholes model from observed option prices.

In-the-Money A situation where an option has intrinsic value, such as when a call option has an exercise price less than the asset price or where a put option has an exercise price greater than the asset price.

Inflation A sustained increase in the general level of prices for goods and services.

Inflation Risk The risk of economic loss due to an increase in inflation.

Initial Margin The initial performance guarantee deposit required by a futures exchange at the inception of a trade.

Insurance Risk The risk of loss associated with the failure of an insurance company to pay out on a policy. Wrapped instruments, including municipal bonds and asset-backed securities that carry guarantees issued by a monoline insurance company, face this risk.

- Interest-Only Strip** An instrument that pays only interest, such as the by-product of creating a zero-coupon bond or a tranche of a collateralized mortgage obligation.
- Interest Rate Cap** A derivative contract that pays the holder the greater of the difference between the reference interest rate and the cap rate and zero. Effectively a series of interest rate call options—caplets—with a strike price equal to the cap rate, it protects the holder against an increase in floating rates.
- Interest Rate Collar** A combination of a long position in an interest rate cap and a short position in an interest rate floor, it effectively allows the holder to lock a floating rate of interest in a collar, or band.
- Interest Rate Floor** A derivative contract that pays the holder the greater of the difference between the floor rate and the reference interest rate and zero. It provides a payout to the holder when floating interest rates are below the floor, and can be used in combination with a floating rate asset to guarantee an income stream.
- Interest Rate Risk** The risk of loss due to a change in the prevailing level of interest rates. Borrowers who have floating rate liabilities face this risk, as well as fixed-income investors seeking to dispose of a fixed-rate security before its final maturity.
- Interest Rate Swap** A contract providing for the exchange of interest, but not principal, payments on a defined notional principal. In a typical fixed-for-floating interest rate swap, one party, said to be holding the (pay) *fixed leg*, pays a fixed rate of interest on the notional and receives payments, calculated using a floating rate index, such as LIBOR. The other party, holding the (pay) *floating leg*, receives fixed rate payments from the first and makes floating payments to it. Interest rate swaps can also involve the exchange of two sets of floating rate payments—see *Basis Swap*.
- IO** See *Interest-Only Strip*.
- Jensen's Alpha** Return on an asset in excess of the risk-free rate r_f and uncorrelated with a broad market return less r_f : $\alpha_{\text{Jensen}} = (r_{\text{asset}} - r_f) - \beta(r_{\text{market}} - r_f)$. Also known as *abnormal return*.
- Jump-Diffusion Process** A model of price movements that combines continuous price diffusion with abrupt discontinuous jumps.
- Key-Rate Duration** The sensitivity of a bond or portfolio to a change in one of a set of interest rates, while holding the others constant.
- Knock-In Option** See *Barrier Option*.
- Knock-Out Option** See *Barrier Option*.
- Lambda** See *Omega*.
- Legal Risk** The risk of loss due to the imposition of legal penalties and damages.
- Legislative Risk** The risk of loss due to the enactment of laws that adversely affect the conduct of a business or industry, such as changes in environmental or tax law.
- Letter of Credit** A document issued by a bank that guarantees an obligation up to a fixed amount. In effect, it replaces the credit risk of the holder of the letter with that of the bank that issued the letter.

Lévy Process A stochastic process governed by probability distributions with algebraically long tails such that the variance of a step size is infinite. Sometimes propounded as a more realistic alternative to the commonly used normal description of logarithmic returns.

LIBOR London Interbank Offered Rate, the average rate at which major banks are willing to lend, usually in Eurodollars. Fixed every day shortly before 11 am London time by the British Bankers' Association.

Linear Instrument A derivative instrument whose price change is approximately proportional to a change in the price of the underlying; for example, a futures contract.

Linear Regression A statistical method of identifying the relationship, if any, that exists between two variables.

Liquidity Risk The risk that the availability of willing buyers or sellers of a given security in the market will be limited, or nonexistent. Limited liquidity typically results in a comparatively lengthy delay before a trade can be fully executed and significant price impact, meaning that if a transaction is completed, it will not be at the theoretical price or fair value. Liquidity risk can also refer to the risk that an entity will be unable to honor an obligation, due to a cash shortfall and lack of other liquid assets and credit facilities.

Listed Option A put or call option traded on an exchange. See *Over-the-Counter*.

Log-Normal Distribution The probability distribution, obeyed by a random variable whose logarithm is distributed normally. Logarithmic (continuously compounded) returns of an asset with a log-normal price are normally distributed and, therefore, are amenable to analytical treatment and numerical modeling. This explains the preponderance of the log-normal price assumption in the analysis of capital markets.

Long Hedge A hedge consisting of a long futures position.

Long Position A position held by the purchaser of an asset. The party that holds the floating, that is, "receive fixed," leg of a fixed-for-floating interest rate swap.

Longevity Risk The risk that purchasers of instruments such as annuities will live longer than anticipated or than was modeled using mortality tables.

Look-Back Option A path-dependent option whose payoff depends on the minimum or maximum asset price that obtained during a specified period. The payout on a *fixed-strike* call is $\max(P_{\max} - K, 0)$ where $P_{\max}$ is the maximum price achieved by the underlying during the life of the option; the payout on a fixed-strike put is $\max(K - P_{\min}, 0)$ with $P_{\min}$ being the minimum price achieved. For *floating-strike* call and put the payouts are $P_T - P_{\min}$ and $P_{\max} - P_T$ respectively, that is, the strike is set to $P_{\min}$ for the call and $P_{\max}$ for the put. (Note that the floating-strike payouts can be zero but cannot become negative.)

Loss Given Default A conditional risk measure, defined as the loss that would be suffered were an obligor to default.

LYON A Liquid-Yield Option Note is a zero-coupon convertible bond that is callable by the issuer and puttable by the purchaser.

Macaulay Duration The weighted average time to the receipt of the cashflows of a fixed-income instrument such as a bond. See *Duration*.

- Maintenance Margin** The minimum permissible balance in a trader's margin account. If the balance falls below the maintenance margin, the trader will receive a margin call and must deposit additional funds.
- Manager Risk** The risk of fund underperformance relative to a benchmark due to a manager having poor selection return or not following the investment objectives of the fund.
- Margin Call** A demand for extra margin to be deposited due to the balance in the margin account falling below the maintenance margin.
- Mark-to-Market** The process of revaluing an instrument to reflect current market conditions.
- Market Exposure** See *Market Risk*.
- Market Impact Cost** The impact on the market price of an instrument attributable to an investor's purchase or sale transaction. When liquidity is high, and markets are deep, market impact cost is low. When the size of a transaction is a significant proportion of, or exceeds, average daily volume in the instrument, market impact costs are likely to be relatively high.
- Market Maker** A broker-dealer that quotes both bid and ask prices and is prepared to buy or sell a specified quantity of a security.
- Market Risk** The risk of loss as the result of an adverse move in the value of a portfolio in response to changes in market factors such as stock and commodity prices, interest and exchange rates, volatility, etc.
- Market Risk Premium** The additional return over the risk-free rate required by an investor to compensate for his assuming the systematic risk of the market.
- Market Sentiment Risk** The risk of loss arising from a complete absence of market liquidity due to a market crash or a prevailing sense of fear and uncertainty in the market. See *liquidity risk*, which relates to a lack of liquidity for an individual security. Market sentiment risk refers to a lack of liquidity for all securities.
- Matrix Price** An estimated price for a fixed-income security, based on the observed prices for other instruments that share the same coupon and duration, or interpolated from the prices of closely related securities.
- MBS** See *Mortgage-Backed Security*.
- Median Tail Loss (MTL)** The median size of the loss expected to occur in the tail of a distribution. See *Expected Tail Loss*.
- Middle Office** The functional areas in a financial organization responsible for supervision, risk oversight and reporting, and profit and loss accounting. See *Front Office* and *Back Office*.
- Model Risk** The risk arising from the use of a quantitative model to value securities. A model is only an approximation of reality, and errors can arise from poor model specification and implementation and from faulty inputs.
- Modified Duration** A measure of the sensitivity of a bond or related instrument to a change in interest rates. It is a variant of Macaulay's original formula for *duration* (*q.v.*). "Duration" nowadays generally refers to "modified duration." Modified duration is usually positive, although instruments such as IOs can have negative duration.

- Moment** For a random variable the moment of order ν is the expected value of this variable raised to the power of ν (usually, but not necessarily, integer).
- Monetary Policy** The control of the supply of money and its cost—the level of interest rates—by a central bank or other governmental authority as a means of effecting policy to influence the level of macroeconomic activity and inflation in an economy.
- Moneyness** The degree to which an option is in-the-money.
- Monoline Insurer** An insurance company that guarantees (“wraps”) bonds in return for a premium. Bonds wrapped by a monoline reflect the rating of the insurance company, except when the rating of the underlying bond is higher than that of the monoline, a phenomenon that occurred for certain bonds during the credit crisis of 2008–2010 when several monoline insurers were downgraded by the rating agencies.
- Monte Carlo Method** A method of asset valuation by stochastic extrapolation. The asset is valued under multiple (typically thousands or tens of thousands) hypothetical market evolution paths and its present value, averaged over the paths, serves as the current value estimate. The method produces a distribution of asset values and thus allows in principle the calculation of all statistics. Distinct types are *historical* and *analytical* Monte Carlo. The former uses either actual realized historical returns or returns generated according to a historically derived distribution. Analytical Monte Carlo chooses returns from a selected distribution or distributions, usually normal for logarithmic returns and Poisson for price jumps.
- Mortgage-Backed Security (MBS)** An asset-backed security where the collateral consists of one or more pools of mortgage loans. The cash flows from the collateral may pass unaltered to the bondholder—a mortgage pass through—or may be tranching, via a waterfall, into a Collateralized Mortgage Obligation, composed of classes with different risk, coupons and expected final maturities.
- MTM** See *Mark-to-Market*.
- Multicollinearity** A situation in a regression analysis where the supposedly independent variables are in fact correlated, leading to regression results that can be unstable and misleading.
- Naked Call** A call option written, or sold, by an investor who does not own the underlying asset.
- Natural Hedge** A position or exposure that reduces or eliminates the risk of some other position or exposure.
- Nominal Rate** The stated level of an interest rate. When an interest rate is paid more frequently than annually, the periodic rate is equal to the nominal rate divided by the number of compounding periods per year. See *Real Rate*.
- Nonlinear Instrument** A derivative instrument whose price changes as a nonlinear function of the change in price of the underlying. For example, a call or put option. See *Linear Instrument*.
- Nonsystematic Risk** The risk unique to a specific firm. It can be eliminated through diversification. Also known as idiosyncratic risk or specific risk. See *Systematic Risk*.

- Normal Distribution** A probability distribution whose density function is an exponential of a second-degree polynomial of random variables. Its widespread use results from its role as the limiting distribution in the central limit theorem and the relative ease of analytical treatment. In quantitative finance, logarithmic returns are often assumed to follow a normal distribution, implying that prices are *log-normally* distributed. Some early students of stock price movements (e.g., Louie Bachelier, Harry Markowitz) assumed, for mathematical convenience, a normal distribution of *prices*, while being aware of the limitations of this approach that admits, albeit typically with a vanishingly small probability, negative prices.
- Notional** The hypothetical principal used to determine payment amounts in derivative contracts such as interest rate swaps. In most types of swaps, the notional principal is not exchanged.
- Notional Limit** The maximum size of a position that a trader or speculator is allowed to enter. Such limits are often used as a form of risk control in firms, particularly for derivative contracts such as swaps.
- Off-Market Swap** A swap whose present value at inception is non-zero. Parties to such a swap will generally exchange a premium at the inception or over the life of the swap. Note that the payment of an upfront premium does not necessarily imply that a swap is off-market; in a prepaid interest rate swap, for example, the present value of the fixed-rate payments can be paid at inception, while present value of the swap is zero.
- Off-the-Run** Any Treasury security other than the most recent issue for a given maturity. Off-the-runs usually carry a higher yield than the most recent issue—the *on-the-run*—to compensate for their lower liquidity. See also *On-the-Run*.
- Omega** The omega, or elasticity, of a derivative is its fractional change in value for a fractional change in the price of the underlying.
- On-the-Run** The most recently issued Treasury security for a given maturity, generally considered desirable because of its higher liquidity. See also *Off-the-Run*.
- Open Interest** The number of futures or option contracts outstanding for a given maturity.
- Open Market Operations** The buying and selling of fixed-income instruments, usually government bonds, by a central bank in order to control the supply of money, and its cost, as measured by interest rates, in an economy.
- Operational Risk** The risk of loss arising from fraud or employee error, or other failures of a firm's operations such as technological or control breakdowns.
- Option** A derivative that gives the owner the right, but not the obligation, to buy or sell an asset at specified price. See *American Option* and *European Option*.
- OTC** See *Over-the-Counter*.
- Out-of-the-Money Option** A call option whose strike price exceeds the spot price, or a put option whose strike price is less than the spot price, of the underlying.
- Over-the-Counter** A transaction that is negotiated between two parties, instead of being conducted on an exchange. Instruments traded over the counter can be highly customized, in contrast to standardized futures contracts and listed options.

- Path-Dependent Option** An option whose payout depends on not just the terminal value of the underlying asset but also the price path it followed. Barrier options and look-back options are path-dependent.
- Performance Attribution** The identification of the factors that led to the change in value of an asset or portfolio. For example, an equity manager's return may be decomposed into style return and selection return, and attributed to the overweighting and underweighting of sectors and the selection or avoidance of individual stocks.
- Plain-Vanilla Bond** A bond that cannot be called by the issuer or put by the purchaser, is not convertible, has a fixed coupon, and pays principal at maturity.
- PO** See *Principal-Only Strip*.
- Political Risk** The risk of loss due to non-market factors such as civil unrest, political instability and unfavorable political decisions such as nationalization and expropriation.
- Portfolio Insurance** A strategy that seeks to reduce or eliminate the market risk of an equity portfolio through the use of a short index futures position.
- Present Value of a Basis Point** See *DV01*.
- Prepayment Risk** The risk to a lender or an investor in a mortgage-backed security that borrowers will repay their loans prior to their final maturity. Usually the result of a decline in interest rates, increased prepayment behavior by mortgagors results in an investor facing reinvestment risk.
- Prepayment Speed** A measure of the rate at which mortgage borrowers in a given pool of obligors are prepaying their loans. A commonly used methodology is the Public Securities Association (PSA) model, with 100% PSA referring to a prepayment rate that increases by 0.2% each month from the inception of a loan through to month 30, after which it remains at a constant rate of 6% per month.
- Price Swap** An arrangement in which one party pays another the price appreciation of a specified reference asset in exchange for a fixed or floating cash flow. Most commonly, the reference asset is a stock or an equity index.
- Primary Dealer** A financial institution that may trade directly with the U.S. Federal Reserve System.
- Principal-Only Strip** An instrument that pays only principal. It can be a by-product of a zero-coupon bond or a constituent of a collateralized mortgage obligation.
- Principal Protected Note** A fixed-income security that attempts to generate an equity-linked return while preserving the investor's principal, which is usually invested in zero-coupon securities.
- Proxy Hedge** A hedging strategy where a derivative on one asset is used to hedge another, highly correlated asset for which no market in derivatives exists. An example would be hedging exposure in one currency using derivatives on another, highly correlated, currency. In this case, one currency acts as a proxy for another. See *Cross-Hedge*.
- Prudent Expert Rule** A requirement that fiduciaries, that is, those with responsibility for the management of investment assets on behalf of a third party, exercise the same duties of care and prudence as a prudent expert would in the conduct of his own affairs.

- Put Option** An instrument that conveys the right, but not the obligation, to sell an underlying asset at a given price on or before a given date.
- Puttable Bond** A bond that can be sold back by the holder to the issuer, at a fixed price, at specified times.
- Puttable Swap** A swap that can be terminated early by one of the parties.
- PV01** The present value change of a portfolio or asset in response to a one basis point move—usually a decline—in interest rates. See also *DV01*.
- Real Rate** The effective yield available on a loan or instrument after accounting for inflation. See *Nominal Rate*.
- Redemption Risk** The risk of loss arising from the departure of investors from a mutual fund or investment pool. The risk is faced by the asset manager, who faces a reduction in management fees, as well as the remaining shareholders who stand to lose if the manager had to sell illiquid assets to fund the redemption.
- Reference Rate** An interest rate benchmark, such as one-month LIBOR, used in the calculation of interest payments in instruments such as floating-rate mortgages, credit cards, and interest-rate swaps.
- Regulatory Risk** The risk of loss arising from unfavorable actions or changes introduced by a regulatory body, such as a change in margin requirements.
- Reinvestment Risk** The risk that the holder of a coupon-paying bond will be unable to reinvest periodic interest payments at a rate equal to or exceeding the yield to maturity (YTM), calculated at purchase. As a result, the achieved yield over the lifetime of the bond may be lower than the one expected by the buyer.
- Relative Delta** Relates a fractional change in the value of an underlying to the dollar change in value of a derivative. See *Absolute Delta*.
- Relative Value Spread** The spread between two related petroleum products, such as heating oil and gasoline.
- Repo** See *Repurchase Agreement*.
- Repurchase Agreement** A sale and repurchase agreement, usually for a short term, whereby a borrower raises funds by selling a security at a haircut while simultaneously agreeing to repurchase it in the future, effectively repaying the loan. The security functions as collateral for the lender.
- Response Matrix** A grid in two or more dimensions K containing asset valuations (or changes therein) for a range of values of K market and non-market factors.
- Return over Maximum Drawdown** See *RoMaD*.
- Rho** The sensitivity of the price of an option or other derivative to a change in interest rates.
- Risk-Adjusted Return** A measure of an investment's return that takes into account its riskiness. See *Sharpe Ratio*.
- Risk Reversal** See *Zero-Cost Collar*.
- Risk Triad** The three elements of a risk management and hedging system. They are the *identification*, *quantification*, and *management* of risk exposures.
- RoMaD** The annualized return over the period divided by the maximum drawdown.

R-Squared See *Coefficient of Determination*.

Scenario Grid A response matrix for coordinated simultaneous market moves.

Selection Return The component of a portfolio's return attributable to the individual securities in which the manager invested, as opposed to the return deriving from market sector allocation.

Semideviation The square root of the semivariance.

Semivariance Unlike variance, which is a measure of the average deviation of a set of observations, such as returns, from their mean, semivariance considers only those observations that are below the mean. By excluding those returns in excess of the mean, attention is paid only to negative deviations, those that are most likely to be of concern to investors. When a target return is used in place of the mean, the measure is referred to as *target semivariance*.

Settlement The process of completing a transaction. Most futures contracts can be settled for cash, where the final payment exchanged is the gain or loss. Commodity futures can also be physically settled, whereby the specified quantity of the physical commodity is exchanged at a specified location such as a designated warehouse. Swap contracts usually have periodic net settlement, where only the difference between the payments to be made by each counterparty are exchanged.

Shadow Delta See *Effective Delta*.

Sharpe Ratio A risk-adjusted return measure that is defined as the return in excess of the risk-free rate, divided by the standard deviation of returns. A disadvantage of the measure is that it penalizes positive, as well as negative, deviations from the mean. In addition, it implicitly assumes that the distribution of returns is entirely defined by its mean and standard deviation, i.e., is normal.

Short Hedge A hedge consisting of a short futures position.

Short Position A position in which a trader has sold more of an asset or instrument than he has purchased, or in which an option writer does not own the underlying asset. The pay-fixed leg of an interest rate swap.

Short Sale The sale of borrowed securities, in an attempt to profit from an expected decline in their price when the security is repurchased in order to be returned to the lender. The borrower deposits cash collateral with the lender of the securities, and the lender keeps some or all of the interest earned on the collateral as the lending fee. The proportion of interest returned is referred to as the rebate rate.

Short Squeeze A situation where a stock price rises rapidly and short sellers rush to repurchase the shares they have sold short, further driving up the price.

Sortino Ratio A risk-adjusted return measure that measures return relative to a benchmark or target, divided by the standard deviation of the target returns.

Sovereign Risk The risk of loss due to the repudiation of its obligations by a sovereign nation, or due to the imposition of exchange controls. See *Political Risk*.

Specific Risk See *Nonsystematic Risk*.

Speed A risk sensitivity of nonlinear assets and portfolios, it is the sensitivity of the gamma to changes in price of the underlying.

- Spot Market** A market where assets are purchased and sold for immediate delivery, as opposed to a futures market, in which delivery is deferred.
- Spot Rate Curve** A curve of yield on hypothetical zero-coupon government bonds.
- Spread Risk** The risk of loss due to a widening of credit spreads, usually as the result of a perceived increase in the credit risk of an industry or a specific issuer.
- Stack-and-Roll Hedge** A hedging strategy where an entire economic exposure is hedged using short-dated futures contracts. As those contracts approach maturity, the hedge is rolled out to the next nearest contract month. See *Strip Hedge*.
- Standard Deviation** A measure of the dispersion of a set of data, equal to the square root of the variance.
- Statement of Financial Accounting Standard (SFAS)** See *Financial Accounting Standard*.
- Static Hedge** A hedge that, once entered, does not need to be adjusted.
- Static Pool Analysis** An examination of a pool of loans to identify their prepayment and default characteristics. The information gleaned from the analysis is often used to make projections about the behavior of borrowers in subsequent pools.
- Stress Testing** The study of a portfolio's response to large adverse market moves.
- Strike Price** See *Exercise Price*.
- Strip Hedge** A hedging strategy where an economic exposure is hedged using futures contracts that expire in the month corresponding to the exposure, in contrast to a stack-and-roll hedge where exposures of varying maturities are hedged with short-dated contracts that are continually rolled over.
- STRIPS** Separate Trading of Registered Interest and Principal of Securities. Zero-coupon bonds issued by the U.S. Treasury and created by separating ("stripping") a Treasury note or bond's coupon and principal components. The Treasury can create as well as retire STRIPS. See *Zero-Coupon Bond*.
- Structured Note** A general term for a security whose payout depends on the performance of one or more reference instruments (indices, securities, pools of collateral, etc.) and on the rules defining the cash flow waterfall. In addition to the performance of the reference instruments, the investor is exposed to the issuer's credit risk.
- Structured Product** A security whose payout is determined by the performance of designated collateral, as well as the rules defining the cash flow waterfall.
- Structure Risk** The risk that an investor may suffer an unexpected loss because he is unfamiliar with the way in which a structured product is designed or with the provisions of the cash flow waterfall.
- Style Analysis** A method used to identify the risk and return characteristics of an investment portfolio. *Holdings-based* style analysis analyzes constituent securities of a portfolio to determine its overall character. *Returns-based* style analysis is a technique developed by William Sharpe, in which the returns of a portfolio are compared to those of a series of benchmark indices using a regression analysis.
- Style Return** That portion of the return of a portfolio attributable to the returns of the market sectors in which it is invested. See *Selection Return*.

Subprime Mortgage A mortgage loan where the credit risk is significantly higher than that of a prime mortgage. A subprime mortgagor is likely to have a poor credit record, few savings, high debt-to-income ratio, and a poor repayment history.

Swap A derivative contract in which one counterparty—or leg—exchanges periodic cash flows with another. Swaps usually feature the regular exchange of payments, have a finite maturity, are over-the-counter contracts, have a notional principal but often do not involve an exchange of principal, and usually have a net present value of zero at inception, meaning that the initial expected economic value of both legs is equal. To avoid potential confusion, parties to a swap will explicitly confirm, at inception, the payment streams they expect to make or receive. For example, one party to a fixed-for-floating interest rate swap will state that he undertakes to “pay fixed and receive floating,” while the other party will agree to “pay floating and receive fixed.”

Synthetic CDO A form of collateralized debt obligation where the underlying credit exposure is via a basket of credit default swaps instead of a portfolio of bonds or other credit obligations. Each class, or tranche, of the CDO will receive cash flows that originate as premium payments on the credit default swaps, in contrast to the cash flows of a traditional CDO, which arise from coupon or interest payments on credit obligations.

Systematic Risk The risk to which an investor is exposed by being invested in the market. The systematic risk of an individual security is the beta of its returns relative to those of the market.

Tail Risk Measures Measures of risk associated with rare events. Examples include value at risk (VaR), expected tail loss (ETL), and median tail loss (MTL). Tail risk measures explore the characteristics of rare events, in the tail of a distribution, in contrast to observations clustered around the mean.

Target Semivariance See *Semivariance*.

Theta The change in the value of an asset resulting from the passage of time. This change is usually, but not always, negative. An example of a security with a negative theta is a European put option deep in the money.

Time Decay A decrease in the value of a financial derivative such as an option as the exercise date approaches. See *Theta*.

Total Return The aggregate return earned over a period of time, consisting of both change in price—price return—and distributions of interest and dividends—income—paid to the investor.

Total Return Swap (TRS) A swap in which one party pays the total return of an asset, including income and capital gains or losses, and receives a fixed or floating cash flow. When the asset is an equity or an equity index, the swap is also referred to as an equity swap. For example, one party—the equity leg—could pay the return on the S&P 500 index, in exchange for receiving from the counterparty—the floating leg—the return on one-month LIBOR less (sometimes plus) a spread.

Tracking Error The standard deviation of the excess returns of a portfolio over its benchmark.

- Trading Security** A debt or equity asset that is being held with the intention of being sold within a short period time. It is reported at fair value with any gain or loss showing on the income statement.
- Tranche** A class or portion of a deal, especially an asset-backed security. Cash flows are allocated among the various tranches according to rules expressed in a waterfall.
- Transition Matrix** A matrix containing the probabilities that bonds or other credit obligations with one credit rating will migrate to another rating category within a specific period of time, usually one year.
- Treynor Ratio** A portfolio's excess return over a benchmark, divided by the portfolio's beta.
- Ultima** Also known as DvolgaDvol, it is the sensitivity of volga to a change in volatility of the underlying.
- Unfunded Hedge** A hedge position consisting of derivatives that have no intrinsic value at inception, such as futures contracts. Such a hedge does not require the deployment of capital, other than that required as margin. See *Funded Hedge*.
- Up and In** A type of knock-in barrier option, one that is activated by the underlying asset price rising through the barrier price.
- Up and Out** A type of knock-out barrier option, one that ceases to exist after the underlying asset price rises through the barrier price.
- Upsilon** Related to the vega, the *upsilon* is the sensitivity of the price of an instrument (usually a variance swap or option) to the variance of returns of the underlying.
- Utility Function** A measure of the benefit or detriment afforded by a gain or loss on an investment. Often used to determine the degree of risk aversion of an investor.
- Valuation Risk** The risk associated with the process used to value prices, for example, that used by a mutual fund company to calculate a net asset value for a fund. If the fund is mispriced, investors will be buying or selling shares at prices that are either too high or too low.
- Value at Risk (VaR)** The maximum loss expected to be suffered by an enterprise in a stated period of time, with a stated degree of confidence. This measure does not provide any information about the distribution of losses in excess of the VaR.
- Vanna** The sensitivity of a derivative's vega to changes in the price of the underlying, or the sensitivity of its delta to changes in its volatility.
- VaR** See *Value at Risk*.
- Variance** A measure of the typical dispersion of a set of points around their mean, or average, value.
- Variance Swap** A swap where one leg pays the actual realized variance for a period in return for a fixed variance.
- Variance Vega** See *Upsilon*.
- Variation Margin** The additional amount that must be paid in response to a margin call, to bring the margin balance up to the minimum (or maintenance) margin requirement.

Vega The sensitivity of the price of an instrument, usually an option, to changes in the standard deviation of the price returns of the underlying.

Vega Convexity See *Volga*.

Volatility Annualized standard deviation of logarithmic returns.

Volatility Gamma See *Volga*.

Volatility Ratio The ratio of the volatility of a portfolio to the volatility of its benchmark.

Volatility Risk The risk of loss due to a change in volatility.

Volatility Surface A graph or table showing implied volatilities for different strikes and maturities.

Volatility Swap A swap on an asset's volatility. Cannot be replicated and is, therefore, difficult to hedge. Has been mostly supplanted by *variance swap*.

Volga The second derivative of the value of a asset, such as an option, with respect to the standard deviation of the price returns of the underlying. It is also known as vega convexity, vomma, and *volatility gamma*.

Vomma See *Volga*.

Waterfall The set of rules defining the order with which available cash should be distributed among various classes or tranches in a structured product such as a collateralized mortgage obligation.

Window Dressing The manipulation of financial statements or portfolio holdings to imply that an entity's financial position, risk profile, or performance is better than it actually is.

Yield to Maturity (YTM) The yield, achieved on an investment in a bond, if the bond is held to maturity and its periodic interest payments are reinvested at the same yield. Equivalently, the single discount rate that, when applied to all of a bond's cash flows, produces its market price.

Zero-Cost Collar A collar in which the strikes of, for example, a short call and a long put are set such that the sale proceeds of the call are sufficient to pay the premium for the put

Zero Coupon Bond A bond, issued at a discount, that does not pay a coupon. Interest accretes over time, with the bond being redeemed at par at maturity. See *STRIPS*.

Zomma Also known as $D\gamma D\sigma$, it measures the sensitivity of the gamma to changes in volatility, or the vanna's sensitivity to price.

Index

- 12b-1 fees, 28, 259
- 1987 stock market crash, 211
- 2007–2010 crisis, ix, 31, 41, 55, 62
- ABS, 23, 183, 259
- ABX index, 220
- Accounting
 - GAAP, 191
 - hedge, 191
- Adjustable rate mortgages (ARMs), 63
- ADR, 27, 259
- Adverse selection, 39
- Algorithmic trading systems, 52
- Alignment of interests, 46
- Alpha, 259
 - Jensen's, 269
- Altman's Z-score, 217
- Amaranth Advisors, 178
- American depository receipt (ADR), 27, 259
- American Insurance Group (AIG), 219
- Annualization, 157, 247–248
 - factor, 151, 158, 248
- Annuity, 180
- Approvals, 189
- Arbitrage, 83–84, 259
- Asset-backed security (ABS), 23, 183
- Asset
 - extrinsic value, 41
 - factor response, 70
 - intrinsic value, 1–2, 41, 54, 70
 - model, 51, 110
 - price, 18, 51, 54, 59
 - risks, 56
 - risky, 83–85, 97
 - valuation, 70
 - value, 51–52, 54, 56, 58–59
- At-issue criteria, 28–29
- Bachelier, Louis, 273
- Back-testing, 189, 260
- Backwardation, 198, 260
- Barclays Capital U.S. Government/Corporate Index, 211, 219
- Bars, 38
- Basis, 48, 194, 260
- Basis point, 260
 - dollar value of a, 21, 132, 265
 - present value of a, 132, 274
- Basis risk, 25, 48, 157, 208, 211, 222, 260
- Behavioral finance, 11, 56, 171, 260
- Benchmark, 35, 211, 260
- Beta, 110, 154–155, 158–159, 261
- Beta adjustment, 213
- Black box, 30, 104, 149, 261
- Black-Scholes formulae, 70
- Black-Scholes-Merton framework, 73,
 - 83–84, 100, 117
- Black-Scholes model, 201
- BofA Merrill Lynch High Yield Constrained Index, 219
- Bogey, 36, 261
- Bond
 - amortizing, 23, 259
 - asset-backed, 23
 - bullet, 23, 38, 82, 130, 132, 134–135, 261
 - callable, 25–26, 130, 133, 261
 - convertible, 24, 130, 213, 262
 - coupon bearing, 24
 - credit spread, 24
 - due diligence, 28
 - floating rate, 117
 - high grade, 130, 132
 - interest-only, 21
 - mortgage-backed, 27
 - plain vanilla, 24–26, 31, 274
 - puttable, 24, 130, 275
 - sinking fund, 23
 - zero-coupon, 21, 24, 71–72, 74, 128–129, 280
- Bootstrapping, 76, 261
- Boundary conditions, 87

- Breaking the buck, 261
- British Bankers' Association (BBA), 77
- Bubble
 - dot-com, 2, 41
 - market, 20
- Bureau of Labor Statistics, 64
- Burnout, 221
- Business cycle, 2, 9–12
 - initial recovery, 9
 - late expansion, 9
 - recession, 9
 - slowdown, 9
- Busted convertible, 24, 261
- Calendar risk, 196
- Call
 - margin, 270
 - option, 33, 200, 261
 - covered, 33, 263
 - delta, 202
 - hedge ratio, 205
 - naked, 33, 272
 - on enterprise upside, 215
 - prepayment as, 62
 - volatility risk, 34
 - provision, 25
 - risk, 25, 44
- Cap, 261
- Capital asset pricing model (CAPM), 154, 250
- Capital gain, 19
- Caplet, 261
- Captive finance company, 44
- Cash collateral account, 29
- Cash flow
 - deterministic, 81
 - projecting, 80
 - waterfall, 28, 81, 261
- CBOE S&P 500 Implied Correlation Index, 146, 148
- CDO, 31, 38, 183, 261
 - synthetic, 38, 278
- CDS, 261
 - cash settlement, 218
 - physical settlement, 218
- CDX index, 220
- Central bank, 6
- Central limit theorem (CLT), 95, 97, 243
 - rapidity of convergence, 244
- Charm, 123, 261
- Cherry picking, 39
- Chevrolet Vega, 110
- Chicago Board Options Exchange (CBOE), 146
- Chicago Mercantile Exchange (CME), 47, 78
- Cholesky decomposition, 243
- Churches, 38
- Citron, Robert, 33
- Claim
 - contingent, 262
- Clearing house, 45
- CLO, 31, 261
- CMO, 27, 261
 - tranches, 27
- Coefficient of determination (R^2), 60, 251–252, 262
- Collar, 262
 - zero cost, 280
- Collateralized debt obligation (CDO), 31, 38, 183, 262
- Collateralized loan obligation (CLO), 31, 262
- Collateralized mortgage obligation (CMO), 27, 262
- Collateral servicing, 28
- Color, 125, 262
- Commodity Futures Trading Commission, 43
- Commodity trading advisor, 43
- Comparative advantage, 175
- Confidence interval, 52
- Consumer price index, 5
- Contango, 199, 262
- Contingent claim, 87
- Continuous compounding, 262
- Convergence, 126
- Convexity, 117, 132, 152, 262
 - effective, 133
 - negative, 26
 - pass-through, 62–63
 - positive, 22, 26, 133
 - vega, 280
- Correlation, 55, 253–254
 - coefficient of, 98, 158–159, 253–254, 263
 - implied, 92
- Cost
 - market impact, 271
 - of carry, 194, 263
 - to close, 44, 263
- Council of economic advisors, 12

- Counterparty, 185
 - limits, 185
- Covariance, 96, 98, 100, 158, 263
- Covered call writing, 33
- CPI, 5
- Credit default
 - spread, xiii
 - swap (CDS), 40, 218, 263
- Credit
 - derivative, 218
 - enhancement, 28
 - event, 220, 263
 - rating, 185, 216, 263
 - issuer, 216
 - issuer-specific, 216
 - risk, 185
 - spread, 20
 - corporate, 20
- Cross-default provision, 220, 264
- CTA, 43
- Cuba, 39
- Cumulant, 91, 234–235, 264
- Cumulative distribution function (CDF), 90,
93, 165–166, 232
- Currency peg, 182
- Curtailment, 264
- Curve
 - discount, 79, 81, 118, 126
 - zero-coupon, 128
 - Phillips, 5
 - spot rate, 72, 276
 - swap, 24
 - yield, 22, 35, 59
 - risk, 22
- Day count convention, 224
- Decay
 - delta, 264
 - time, 278
- Default
 - risk, 214
 - tables, 218
- Deflation, 5, 264
- Delinquency rate, 109
- Delta, 112–113, 201, 264
 - absolute, 113–114, 259
 - decay, 123
 - effective, 203, 265
 - hedging, 25, 48, 173, 210, 264
 - effective, 202, 204
 - modified, 203
 - neutral portfolio, 201, 264
 - relative, 113–116, 203, 275
 - riding, 204–205, 264
 - shadow, 203, 276
- Derivative, 264
 - credit, 263
 - interest rate, 112
 - mathematical, 111, 141
 - functional, 134
 - logarithmic, 114
 - partial, 112, 118, 131
 - secant, 203
 - total, 131
 - non-linear, 111
 - over-the-counter, 45
 - pricing, 187
 - time decay, 118
 - volatility, 116
- DgammaDvol, 264
- Disaster myopia, 176
- Discount curve, 79, 81, 118, 126
 - flat, 130
 - parallel shift, 118, 126–127, 129, 131, 133
 - sensitivity, 137
 - shape, 133
 - zero-coupon, 128
- Discounted cash flows (DCF), 54, 70–71, 74,
127–128, 265
- Discount rate, 7–8, 265
- Discount window, 7
- Distribution, 254
 - binomial, 254
 - Gaussian, 239
 - leptokurtic, 236
 - Levy, 94–95, 97, 100, 153, 164, 168–169
 - limit, 97
 - log-normal, 87, 94, 239–240
 - multivariate, 241–242
 - normal, 243
 - normal, 90–91, 96–97, 236, 239
 - platykurtic, 236
 - Poisson, 100, 231, 240, 272
 - Student's *t*, 254
 - tails, 89
 - unimodal, 238
- Diversification, 37, 184, 193, 265
- Diversifier, 54
- Dividend income, 19
- Dollar value of a basis point, 21

- Dot-com bubble, 265
- Down-and-in, 265
- Down-and-out, 265
- Drawdown, 155, 157, 162, 265
 - maximum, 157
 - recovery period, 155
 - return over maximum, 275
- Duration, 117, 126–127, 129–130, 132, 152, 265
 - and hedging, 205
 - bond, 20
 - effective, 21, 129–133, 203, 266
 - key-rate, 22, 117, 126, 133, 269
 - Macaulay, 127–129, 270
 - Macaulay-Weil, 128
 - mismatch, 22
 - modified, 20, 127–129, 131–132, 271
 - negative, 21, 29, 131
 - neutrality, 133
 - pass-through, 62–63
 - spread, 21
 - true, 128
- DV01, 21, 132, 224, 265
- DvegaDtime, 124, 265
- DvolgaDvol, 126, 265
- Dynamic hedging, 200, 212
- Early expansion, 9
- Economic indicators, 265
 - ADP employment report, 13
 - car sales, 13
 - consumer confidence, 13
 - durable goods orders, 13
 - EIA petroleum status report, 13
 - employment situation report, 13
 - manufacturing index, 13
 - pending home sales, 13
- Economic report of the president, 12
- Economic value, 51–52, 54
- Efficient markets, 11
- Elasticity, 67–68, 110, 116, 266
- Embedded derivatives, 192
- Endowment funds, 32
- Energy futures, 194
- Enron, 215
- Equivalent variation investment (EVI), 115, 139, 206, 266
 - and dynamic hedging, 210
 - and factor sensitivities, 103
 - and hedge notional, 48
 - and hedge sizing, 160
 - and quantifying risk, xv
 - and relative delta, 115
 - effective, 203, 213
 - uncertainties, 115
- ERISA, 174
- Error
 - amplification, 67–68, 125, 139
 - data flow, 65–66
 - function, 70, 266
 - implementational, 66–61, 70
 - numerical, 67–69, 93
 - operational, 43
 - programming, 66
 - roundoff, 67
 - simulation, 68
 - statistical, 70
 - tracking, 53, 157, 211
 - truncation, 67
- ETF, 25, 33, 43
- Euro STOXX 50 Index, 105–107, 109, 121, 140
- Excess spread, 29
- Exchange-traded funds (ETFs), 25, 33, 43
- Execution, 189
- Expected loss, 217
- Expected tail loss (ETL), 164, 167–169, 266
- Expected value, 85, 233–235, 243
- Expiration price, 187
- Exposure hedge ratio, 186, 206
- Factor return, 92
- Factors of production, 54
- Factor universe, 80–82, 87, 100
 - path, 81–82
 - probability, 81
 - realization, 80–82, 87
- Failed trade, 45
- Fair value, 30
- Fannie Mae (FNMA), 63
- FAS 115, 191
- FAS 133, 191
- FAS 157, 30
- FAS 87, 211
- FASB, 30, 211
- Federal funds, 7
 - rate, 7
 - target, 7
- Federal Home Loan Banks (FHLB), 63

- Federal Open Market Committee, 7, 61, 147, 267
- Federal Reserve Act, 7
- Federal Reserve System, the, 7
- Fed, the, 7
- Financial Accounting Standard (FAS), 267
- Financial Accounting Standards Board, 30, 211
- Financial asset, 19
- Financial crisis of 2007–10, ix, 55, 62
- Financial distress
 - cost of, 174
- Financial instrument, 19
 - price, 1–2, 18, 51, 54, 59
 - value, 1, 10, 18–19, 30, 51
- Fisher, Irving, 4
- Flight to quality, 24, 184, 215
- Floor, 267
- FOMC, 7, 147, 267
- Forced buy, 52
- Forced liquidation, 184
- Forced sale, 52
- Forward rate agreement (FRA), 78
- Foundations
 - charitable, 32
- Fraud, 43
- Freddie Mac (FHLMC), 63
- Free float, 36
- Frequency
 - of compounding, 128
 - of data observations, 64, 98–99, 151
 - of outcomes, 164, 227–229
 - of rebalancing, 202
- Frictional losses, 54, 84, 203, 205
- Fund
 - actively managed, 53
 - exchange-traded, 25, 33, 43
 - index, 53
- Futures
 - contract, 47, 112, 267
 - heating oil, 194
- Gain on sale, 215
- Gamma, 119, 121, 123, 138, 202, 267
 - climbing, 204, 267
 - contribution to effective delta, 203
 - diagonal, 122, 140
 - hedging, 122, 202
 - negative, 49
 - neutral strategy, 202, 206
 - numerical accuracy, 139
 - of an option, 48, 122, 224
 - positive, 122, 210
 - versus time decay, 119
 - volatility, 280
- Gate provision, 35, 268
- Gaussian, 96–97, 100, 164
- GDP, 9, 12
 - deflator, 13
- Ginnie Mae (GNMA), 63
- GNP, 9
- Government-sponsored enterprise (GSE), 62–63
 - charter, 63
- Greater fool theory, 20, 268
- Great Recession, the, ix
- Greeks, the, 110, 268
 - and interest rate sensitivities, 127
 - linear, 112
 - second order, 123
- Gross domestic product, 9
- Gross national product, 9
- Haircut, 268
- Half-life, 64, 98
- Headline risk, 45
- Hedge, 115
 - cross, 182, 264
 - delta, 264
 - equity futures, 206
 - execution, 223
 - funded, 209, 267
 - funds, 212
 - gamma, 267
 - imperfect, 208
 - layering, 213
 - linear, 210, 223
 - long, 270
 - natural, 272
 - non-linear, 47, 210, 223
 - proxy, 206, 221, 274
 - ratio, 186, 205
 - reasons to, 173
 - short, 276
 - sizing, 160, 162
 - optimal, 160, 162
 - stack and roll, 277
 - static, 277
 - strategy, 187
 - strip, 277

- Hedge (*Continued*)
 - tenor, 223
 - theta, 119
 - unfunded, 209, 279
- Hedging, 47, 119, 146
 - dynamic, 265
 - first rule of, 173
 - justification for, 176
 - opportunity cost, 33
 - second rule of, 184
 - third rule of, 207
 - time zone differences, 224
 - unintended consequences, 213
- Historical regression, 64
- Historical test, 93–94
- Hunt brothers, 178
- Index
 - custom liability, 211
 - futures, 182
- Indicators
 - economic, 11
- Inflation, 4, 12, 268
 - and unemployment, 5
- Information asymmetry, 46, 173, 215
- Initial conditions, 87
- Instrument
 - linear, 270
 - nonlinear, 272
 - over-the-counter, 273
- Insurance
 - bond, 261
 - portfolio, 274
- Insurer
 - monoline, 272
- Intercontinental Exchange (ICE), 47
- Interest only security (IO), 21, 29, 131
- Interest rate
 - cap, 34, 269
 - collar, 269
 - floor, 269
 - nominal, 4
 - real, 4
 - term structure, 117, 127
 - volatility, 34, 77, 79, 117
- In-the-money, 268
- Investment
 - objectives, 35
 - portfolio, 87–88
 - risk-free, 25, 72–73, 179
- Investor
 - buy-and-hold, 41
 - psychology, 52
 - short-term, 41
- IO, 21, 29, 269
- ISDA, 186, 220, 263
- iTraxx index, 220
- Jump-diffusion process, 100, 117, 240, 269
- Key maturity, 133–137
- Keynes, John Maynard, 42
- Key rate, 22, 134, 136, 142
 - convexity, 134–135
 - off-diagonal, 135
- decoupling, 135
- deformation, 136
- duration, 134–135
- DV01, 134
- measures, 126, 133–135
- PV01, 134
- KMV Creditor Monitor, 217
- Knockout option, 30, 187
- Kurtosis, 81, 91, 235–236
- Lambda, 116, 269
- Law
 - of demand, 3
 - of large numbers, 11
 - of one price, 201
 - of supply, 3
- Leading indicators, 10
- Lehman Brothers, 72
- Lender of last resort, 7
- Letter of credit, 180, 215, 269
- Leverage, 184
- Levy
 - distribution, 94–95, 97, 100, 153, 164, 168–169
 - flight, 94
 - Paul Pierre, 94
 - process, 94, 100, 269
- Liar loans, 215
- Libor, 64, 77, 270
 - curve, 78
 - swap curve, 78
 - swap rates, 64, 78
- Limit
 - counterparty, 263
 - exposure, 181, 266

- notional, 181, 273
- stop-loss, 181
- Liquidity, 224
 - crisis of 2007–2010, 31
 - premium, 37, 75, 183, 218
- Loan covenant, 173
- Loan-to-value ratio (LTV), 63
- Log-normal distribution, 87, 94
- Long Term Capital Management, 184
- Loss
 - given default, 23, 218, 270
 - given default (LGD), 217
 - mitigation, 215
 - severity, 23
- LYON, 24, 270
- Lyra, constellation of, 110
- Macaulay, Frederick R., 128
- Managed futures fund, 43
- Margin, 185
 - call, 32, 197, 224, 271
 - initial, 196, 268
 - maintenance, 197, 271
 - variation, 197, 279
- Market
 - clearing price, 3
 - discount, 2
 - exposure, ix, xvi, 49, 103, 110, 115, 271
 - factor, 204
 - co-movement, 109
 - sensitivity, 55, 149
 - impact, 31
 - cost, 32
 - maker, 271
 - mark- to-, 271
 - premium, 2
 - sentiment, 99
 - spot, 277
- Markowitz, Harry, 37, 273
- Mark-to-market, 30, 271
 - timing differences, 224
- Matrix
 - pricing, 1, 271
 - response, 275
- MBS, 25, 27, 31, 45, 271
- Mean, 81, 235, 237
- Measures
 - central tendency, 237
 - tail risk, 278
- Median, 81, 238
- Median tail loss (MTL), 164, 167–169, 271
- Merton model, 217
- Metallgesellschaft, 197
- Mispricing, 52
- Mode, 238–239
- Model, 51
 - accuracy, 59, 69, 86
 - numerical, 69
 - saturation, 93
 - behavioral considerations, 56
 - Black-Scholes, 261
 - capital asset pricing, 261
 - complete, 142
 - concept, 56
 - data checkers, 61
 - error amplification, 67, 125
 - factor, 56, 59, 86, 103, 108, 149
 - complete set, 60
 - continuous, 71, 80, 84–85, 91, 103–104
 - discrete, 71, 80–81, 84–86, 91, 104
 - economic, 56, 59, 62
 - external, 64, 70
 - instrument, 115
 - interdependent, 60
 - macroeconomic, 64
 - market, 56, 59, 71, 92, 103–104, 110, 115, 117, 119
 - non-market, 39, 71, 84–85, 91–92, 104, 110
 - qualitative, 57, 64–65
 - quantifying, 56
 - quantitative, 56–57
 - variability, 86
 - financial, 55
 - formalism, 57, 61
 - frugality, 59
 - future projections, 57, 64
 - implementation, 52, 58–61, 65
 - inputs, 56–60
 - delayed, 61
 - subjective, 57, 61
 - political considerations, 56
 - prepayment, 58, 222
 - quantitative, 55–56
 - regressions, 57, 64
 - data frequency, 64, 98–99
 - equal-weighted, 64
 - historical decay, 64
 - historical window, 64

- Model (*Continued*)
- least squares, 64
 - robust, 64
 - risk, 65, 69
 - sensitivity, 69
 - statistical errors, 69
 - test, 58
 - uncertainty, 66, 132, 138–139
 - conceptual, 66, 115
 - formalism, 66, 115
 - implementational, 66
 - numerical, 67
 - unstable, 59
 - valuation, 55–56, 67, 70, 103–104, 132, 137, 139, 149
- Modern portfolio theory (MPT), 37, 171
- Moment, 271
- Monetary policy, 6–7, 271
- Moneyness, 272
- Money
- quantity theory of, 4
 - supply, 4
- Monitoring, 190
- Monoline insurance company, 183
- Monte Carlo method, 64, 88, 91, 94, 166, 272
- analytical, 92
 - asset-backed securities valuation, 88
 - basic theorem, 89
 - derivatives valuation, 70
 - historical, 92–93
 - hybrid, 92
 - limitations, 89, 91
 - random number generator, 68
 - variance reduction techniques, 69
- Moral hazard, 176
- Mortgage
- adjustable-rate (ARM), 63
- Mortgage-backed security (MBS), 21, 25, 60, 130
- conforming loan limit, 63, 65
 - curtailment, 62
 - delinquency rate, 62–63, 109
 - foreclosure rate, 62–63
 - geographic concentration, 38, 62
 - pass-through, 61
 - prepayment, 62
 - model, 65
 - speed, 60, 62–63, 109
 - refinancing, 62
- Mortgage Bankers Association of America (MBAA), 62
- Mortgage
- call provision, 25
 - call risk, 44
 - curtailment, 264
 - fixed-rate, 62
 - Florida, 63
 - hedging with balance guarantee swap, 260
 - liar loans, 215
 - originator, 54
 - pass-through, 27, 65
 - pool, 62, 65
 - prepayment, 28, 221
 - risk, 221, 274
 - speed, 109, 274
 - private insurance, 215
 - savings and loans crisis, role in, 23
 - subprime, 24, 216, 277
 - hedge, 220
 - willful default, 220
- MPT, 171
- MSCI Far East Index, 160
- Multicollinearity, 60, 272
- Naked call writing, 33
- NASDAQ 100 Index, 113–115, 208
- Natural gas, 11
- futures, 178
- Natural hedge, 180
- Nelson-Siegel function, 79
- Net asset value (NAV), 46, 209
- NIKKEI 225 Index, 160, 224
- Normal distribution, 90–91, 96–97, 236, 239
- Notional, 273
- futures, 48
 - option, 48
- Null hypothesis, 253–254
- Offering memorandum, 216
- Office
- back, 260
 - front, 267
 - middle, 271
- Off-the-run, 273
- Oil, 11
- hedging, 182–183, 198, 200
 - basis risk, 213
 - industry nationalization, 39

- market, 177
- price, 19, 42, 44, 175
- reserves, 42
 - proven, 42
 - unproven, 42
- resources, 42
- supply and demand for, 13
- Omega, 116, 273
- One way markets, 188
- On-the-run, 273
- Open interest, 273
- Open market operations, 6–7, 273
- Opportunity cost, 49, 208
- Optimal strategy, 186
- Option, 125, 273
 - American, 87–88, 187, 259
 - Asian, 88, 187, 259
 - at-the-money, 120, 260
 - barrier, 214, 260
 - Bermudan, 88, 187, 261
 - call, 118–120, 200, 261
 - delta, 202
 - hedge ratio, 205
 - naked, 272
 - compound, 262
 - credit spread, 219, 264
 - delta, 202
 - down-and-out, 187
 - European, 70, 87–88, 118, 187, 266
 - exchange-traded, 187
 - exercise style, 187
 - exotic, 266
 - Greeks, 224
 - index, 117
 - interest rate, 113, 117
 - in-the-money, 44, 118
 - knock-in, 269
 - knock-out, 30, 269
 - listed, 45, 270
 - look-back, 88, 188, 270
 - call, 188
 - fixed strike, 188
 - floating strike, 188
 - put, 188
 - naked, 272
 - OTC, 187
 - out-of-the-money, 44, 273
 - path-dependent, 273
 - plain vanilla, 60, 83, 89, 100, 117–118, 122, 201
 - put, 48, 118, 209, 275
 - shout, 188
 - stock, 113, 117–118, 120
 - up-and-in, 279
 - up-and-out, 279
 - variance, 112–113, 116–117
 - volatility risk, 34
- Orange County, 33
- Order
 - stop, 45
 - stop-limit, 45
- Originate to distribute, 216
- OTC, 45, 273
- Outlier, 61, 163, 238
- Out-of-the-money option, 273
- Output gap, 9
- Over-collateralization, 29
- Overhedged portfolio, 196, 208
- Over-the-counter, 273
- Paper trading, 189
- Par coupon curve, 76
- Partial differential equations, 86
- Pass-through, 61
 - convexity, 62–63
 - duration, 62–63
- Path dependence, 87–88, 187
- Pension accounting, 211
- Performance attribution, 56, 139
 - interest rate, 142
 - linear, 140
 - quadratic, 140
 - second order, 142
- Performance triggers, 28
- Perrier, 45
- Personal guarantee, 215
- Phillips Curve, 5
- Places of contemplation, 38
- PO, 29, 274
- Porsche SE, 44, 52
- Portfolio
 - insurance, 210
 - risk-free, 83–84
- Position
 - long, 270
 - short, 276
- Pre-averaging approximation, 82, 85–86
 - partial, 82, 86
- Pre-hedging, 190

- Premium
 - bond insurance, 183
 - credit default swap, 40, 263
 - inflation, 4
 - interest rate cap, 34
 - liquidity, 37, 75, 183, 218
 - market, 2
 - market risk, 271
 - option, 33, 47–48, 122, 186
 - risk, 83, 85, 152, 155, 175, 184, 261
 - swaption, 34
- Prepayment, 21, 29, 62, 130
 - model, 58, 65
 - rate, 64–65, 190
 - speed, 60, 62–63, 66, 88, 109, 274
- Present value, 54, 71–72
- Price
 - ask, 1, 51, 259
 - bid, 1, 51, 261
 - ex ante, 36
 - exercise, 266
 - ex post, 36
 - inefficiency, 183
 - market, 53
 - obscurity, 223
 - observed, 52–53
 - offer, 51
 - settlement, 197
 - stop, 45
 - strike, 277
 - trade, 52
- Primary dealer, 8, 274
- Principal only (PO), 29
- Private mortgage conduit, 62
- Private mortgage insurance (PMI), 215
- Probability, 227–229, 235, 239
 - conditional, 228
 - density, 232, 238–239, 241, 243, 272
 - joint, 241
 - distribution function, 81, 229–231
 - continuous, 232
 - cumulants, 91, 234–235, 264
 - discrete, 231
 - moments, 234, 236
 - a priori*, 228
 - of default, 217
- Process
 - deterministic, 228
 - jump-diffusion, 100, 117, 240, 269
 - Lévy, 94, 100, 269
 - stochastic, 83, 270
- Profit-or-loss (P/L), 48, 139–142
 - report, 190
- Progressive tax scale, 173
- Protective puts, 33
- Proxy hedging, 205
- Prudent expert rule, 174
- Put spread, 30
- PV01, 117, 126, 132, 275
- P*-Value, 146, 253–254
- Quantile, 81
- Random
 - innovation, 86
 - number, 69, 92–93, 96
 - generator, 69, 93
 - phenomena, 227
 - variable, 90, 93, 95–96, 266
 - continuous, 232
 - discrete, 231
- Rate
 - delinquency, 264
 - discount, 74, 79, 117
 - discount (central bank), 7–8
 - floating, 21, 267
 - foreclosure, 267
 - forward, 117, 267
 - agreement, 78
 - nominal, 272
 - of return
 - expected, 117
 - required, 83–84, 117, 172
 - real, 275
 - reference, 74–75, 79, 135–136, 275
 - risk-free, 55, 73–74, 83–84
 - and ρ , 117–118
 - and risky assets, 85, 97
 - and scenario grid, 107
 - and semideviation, 152
 - and Sharpe ratio, 153
 - negative, 72
- Rating
 - agencies, 28
 - through the cycle, 218
- Ratio
 - exposure hedge, 186, 206, 266
 - hedge, 207, 268

- interest coverage, xiii
- price-to-earnings, 46
- Sharpe, 55, 153, 276
- Sortino, 55, 154, 276
- Treynor, 55, 154, 279
- volatility, 158, 160–161, 280
- Refinancing incentive, 221
- Regressand, 248
- Regression
 - linear, 109, 248–249, 270
 - dependent variable, 248
 - independent variable, 248
 - multivariate constrained, 251
 - multivariate unconstrained, 251
 - ordinary least squares (OLS), 249
 - univariate, 250
 - robust, 64, 249
- Regressor, 248
- Regulatory risk, 45
- Replicated portfolio, 206
- Repo, 8, 275
- Repurchase agreement, 8, 275
 - reverse, 8
- Reserve Management Corporation, 73
- Response, 103
 - asymmetric, 26, 122
 - linear, 67
 - matrix, 104–106, 149
 - and stress-testing, 109
 - synchronized, 108–109
 - measure, 104
 - non-additivity, 105
 - non-linearity, 119
- Return
 - absolute, 35
 - alpha, 35
 - arithmetic, 94, 114
 - beta, 35
 - continuously compounded, 94
 - expected, 266
 - idiosyncratic, 160
 - logarithmic, 94, 114
 - and central limit theorem, 95
 - and Monte Carlo method, 94
 - and semivariance, 153
 - and volatility, 116, 151
 - in liquid markets, 96
 - over maximum drawdown (RoMaD), 157
 - price, 114
 - risk-adjusted, 54, 275
 - risk-free, 72, 74, 85, 171
 - risky, 85
 - selection, 35, 276
 - style, 35, 277
 - systematic, 160
 - target, 153
 - total, 278
- Rho, 117–119, 126, 275
 - and hedging, 205
- Richardson extrapolation, 139
- Risk, 44, 145
- Risk-adjusted return, 54
- Risk
 - basis, 25, 48, 157, 211, 260
 - and cross hedge, 182
 - and hedging decision, 49, 115
 - benchmark, 36, 260
 - call, 25
 - commodity, 42, 262
 - compliance, 45, 262
 - concentration, 29
 - conservation of, 177
 - contagion, 27, 146, 262
 - contraction, 62
 - correlation, 37, 263
 - counterparty, 44, 184, 223, 263
 - country, 39, 263
 - credit, xv, 23, 128, 134, 214, 264
 - credit spread, 219
 - definition, 17
 - derivatives, 32, 264
 - diversifier, 54
 - downgrade, 214, 218
 - earnings, 265
 - earnings-at, 265
 - equity, 19, 266
 - exchange, 45, 266
 - exposures, 179, 181–182
 - identifying, 179
 - extension, 28, 62, 266
 - foreign exchange, 27, 73, 267
 - FX, 267
 - headline, 268
 - human behavior, 45, 221, 268
 - idiosyncratic, 154, 268
 - inflation, 37, 73, 268
 - insurance, 27, 268
 - interest rate, 20, 117, 190, 224, 269

Risk (*Continued*)

- key man, 46
- legal, 43, 269
- legislative, 41, 269
- limits, 181
- liquidity, 31, 36, 45, 48, 184, 270
- longevity, 180, 270
- management, ix–x, xvii, 42, 56, 59, 96, 145
- manager, 34, 181, 271
- market, 148, 184, 271
- market sentiment, 31, 41, 271
- measure, ix
 - absolute, 150
 - bulge, 163–164
 - extreme, 155, 162
 - extrinsic, 149
 - intrinsic, 149
 - relative, 149, 157
 - robust, 155, 162
 - standard, 149
 - tail, 150, 162–163
- model, 28, 30, 65, 69, 150, 201, 271
 - sensitivity, 69
- mortality, 180
- nonsystematic, 19, 272
- operational, 43, 221, 273
- people, 43
- political, 39, 56, 63, 73, 274
- portfolio, 44, 145
- potential, 44
- premium, 83, 85, 152, 155, 175, 184, 261
- prepayment, 28, 221, 274
- redemption, 32, 221, 275
- regulatory, 275
- reinvestment, 24, 73, 221, 275
- reporting system, 225
- reversal, 224, 275
- skew, 164
- sovereign, 39, 73, 276
- specific, 19, 276
- spread, 214, 218, 277
- structure, 27, 277
- systematic, 19, 154, 278
- tail, 96–97, 166
 - measures, 91, 150, 162–163
- triad, ix, 103, 275
- uninsurable, 223
- unsystematic, 37, 154
- valuation, 46, 279

- volatility, 34, 213, 280
- willingness and ability to take, 175, 208
- yield curve, 22

RoMaD, 157, 162, 275

Roundoff error, 67

R-squared, 60, 251–252, 275

Rule

- prudent expert, 274

Russell 1000 Index, 26, 66, 69

Russell 2000 Index, 47–48

Russell 3000 Index, 53, 69

Russian debt crisis, 39

Sample history, 212

Savings and loans crisis, 23

Scenario, 64, 104, 108–109

- analysis, 149

- grid, 105, 107–108, 149, 276

- and effective delta, 203

- and stress-testing, 109

- selection, 108

Scholes, Myron, 110

Seagram Company, the, 215

Secondary market, 30

Securities lending, 47

Securitization, 215

Security

- asset-backed, xvi, 23, 27–28, 118, 130, 183, 259

- exotic, 31

- available-for-sale, 260

- hold-to-maturity, 268

- mortgage-backed, 21, 25, 27, 272

- mutual fund fee-backed, 45

- trading maturity, 279

Semideviation, 152, 276

- target, 153

Semi-invariant, 91, 234

Semivariance, 152, 276

- target, 153, 278

Sensitivity, 48, 103–104, 110–113

- and performance attribution, 139

- and Taylor theorem, 112

- cross, 121

- error, 67

- first order, 138, 140

- higher order, 124, 139–140, 204

- interest rate, 117–118, 126–127, 142

- linear, 110–112, 119, 132

- catalog, 112

- numerical evaluation of, 137
- price, 112
- second order, 119, 132, 138, 140
 - catalog, 123
 - symmetry of, 121
- third order, 125
 - catalog, 125
- volatility, 110, 116
- Settlement, 276
- Severity of loss (SOL), 217
- Severity scale, 222
- Shadow gamma, 204
- Sharpe ratio, 55, 153, 276
- Short sale, 32, 115, 276
- Short squeeze, 32, 44, 52, 276
- Skewness, 81, 91, 235
- S&L crisis, 23
- Smallness criterion, 123
- Sortino ratio, 55, 154, 276
- Speculator, 177
- Speed, 125, 138, 203, 276
- Spline, 79
 - cubic, 79
 - higher order, 79
 - quartic, 79
- Split ratings, 217
- S&P 500 Index, 25, 35, 38, 64, 66–67, 69, 105–107, 109, 121, 140, 146–148, 206, 208, 211–212
- S&P 500 Telecommunication Services Index, 113–116
- S&P 500 Total Return Index, 146, 148
- Spot rate curve, 72, 76
 - theoretical, 75
- Spread
 - arbitrage, 259
 - bid-ask, 223, 261
 - bid-offer, 261
 - crack, 177
 - credit, 264
 - credit default, 263
 - relative value, 275
 - risk, 218
- Stack and roll hedge, 198
- Stack hedging, 198
- Standard deviation, 18, 81, 89, 96, 116, 150, 277
- Statement of Financial Accounting Standard (SFAS), 277
- Static hedging, 200
- Static pool analysis, 215, 277
- Statistical significance, 146, 229, 249, 253–254
 - test, 247, 253
- Statistics, 1–2, 11, 229, 247, 249, 253
- Stochastic
 - asset valuation, 82
 - differential equations, 86
 - extrapolation, 100, 150
 - innovations, 87
- Stop price, 45
- Stress-testing, 108–109, 149, 181, 277
- Strip hedge, 198
- Strip
 - interest-only, 269
 - principal-only, 274
- STRIPS, 75–76, 277
- Structured
 - note, 29, 31, 277
 - product, 29–30, 192, 277
- Style analysis, 206–207, 212, 277
- Sub-prime mortgages, 24
- Supply and demand, 3, 19, 52
- Surface
 - volatility, 280
- Svensson function, 79
- Swap, 278
 - amortizing, 259
 - balance guarantee, 190, 260
 - basis, 260
 - basket, 260
 - credit default (CDS), 38, 218
 - cross-currency, 223, 264
 - equity, 266
 - FX, 223
 - interest rate, 44
 - interest-rate, 177
 - interest rate, 269
 - off-market, 223, 273
 - puttable, 275
- Swaption, 34
- Swap
 - total return, 219, 278
 - variance, 116–117, 224
- Tail
 - probability distribution, 168–169
 - risk, 96–97, 162
 - measures, 91, 150, 162–163

- Target
 - semideviation, 153
 - semivariance, 153
- Tax rate, 173
- Taylor
 - expansion, 112, 121, 125
 - rule, 8
 - theorem, 112
- Technical analysis, 11
- Technical default, 220
- Test
 - completeness, 60, 139, 142–143
 - ex post, 59
 - extreme value, 58–59, 66
 - historical, 59, 93–94
 - in-sample, 59
 - out-of-sample, 59
 - self-consistency, 59, 141
 - stress, 108–109, 149, 181
- Theta, 118–119, 278
 - alternative definition, 119
 - and delta riding, 204
 - and hedging, 205, 214, 224
 - break-even move, 119
 - inability to hedge, 119
 - positive, 118
- Time decay, 118–119, 204, 214
- Time value of money, 71–72, 80
- TIPS, 37
- Too big to fail, 176
- Total return, 24
- Tracking error, 53, 157, 211, 278
- Trade ticket, 225
- Tranche, 279
- Transaction costs, 36, 54, 84, 122, 202, 204, 210
- Transition matrix, 217, 279
- Tree
 - binomial, 87
 - multinomial, 87–88
 - non-recombinant, 88
 - recombinant, 88
 - trinomial, 87
- Treynor ratio, 55, 154, 279
- True sale, 216
- Truncation error, 67
- t*-Statistic, 254
- Ultima, 126, 279
- Underhedged portfolio, 196
- Unemployment, 5–7, 12–13, 62–63, 71, 109
 - and inflation, 5
 - rate, 5, 8, 11, 60, 64, 84
- Upsilon, 116, 124, 279
- Utility function, 55, 279
- Valuation
 - risk-neutral, 83–85, 93, 117
 - single-path, 64, 81–82, 85
- Value added, 54
- Value at risk, 18, 91, 164–167, 179, 279
 - confidence level, 164
 - time horizon, 164
- Value
 - expected, 266
 - fair, 267
- Vanna, 124, 279
- VaR, 18, 91, 164, 167–168, 179, 279
- Variance, 91, 150, 235, 279
 - annualized, 248
 - covariance matrix, 96–98
 - of variance, 117
 - vega, 279
- Vega, 110, 116, 123, 213, 280
 - and hedging, 205
 - Chevrolet, 110
 - star, 110
 - variance, 116, 124
- Venezuela, 39
- View of the market, 35
- Volatility, 150, 247, 280
 - and delta riding, 204
 - and response matrix, 105, 109
 - and robust regression, 64
 - and scenario grid, 107
 - and Sharpe ratio, 55
 - effect on asset valuation, 108
 - factor, 82
 - historical, 34, 268
 - implied, 34, 60, 92, 99, 268
 - inflation rate, 6
 - in stochastic extrapolation, 101
 - interest rate, 34, 79, 152
 - map, 203
 - market, 19, 31, 36, 45
 - P/L, 48–49
 - portfolio, 54
 - ratio, 158, 160–161

- risk, 34
- sensitivity, 116
- surface, 214
- swap, 213
- Volga, 124, 280
- Volkswagen, 44, 52
- Vomma, 124, 280
- Waterfall, 280
 - cash flow 261
- Weighted average coupon (WAC), 61, 63
- Weighted average time to maturity (WAM), 62–63
- Window dressing, 73, 280
- Yield curve, 22, 35, 59
 - shape, 22, 133
- Yield to maturity (YTM), 21, 128, 132
- Zero-cost collar, 224
- Zero-sum game, 219
- Zomma, 126, 280